BIG-IP® System: Essentials

Version 11.6

Table of Contents

Table of Contents

General Configuration Properties

About BIG-IP system general configuration properties

Part of managing the BIG-IP® system involves configuring and maintaining a set of global system properties. These properties enable you to configure:

- General device features, such as NTP and DNS
- General local traffic features, including some global persistence settings
- General global traffic features, including load balancing and metric collection

When you configure general device properties, you are affecting the operation of the BIG-IP system as a whole, rather than just one aspect of it. Similarly, when you configure the general properties related to local traffic or global traffic, you are globally affecting the operation of the local traffic management and global traffic management systems.

About general device properties

The BIG-IP® system general device properties that you can view or configure are:

- The host name
- The BIG-IP software version number
- The number of CPUs available
- The number of CPUs that are active

Other BIG-IP system general device properties that you can configure are:

- Network boot
- Quiet boot
- Display of the LCD system menu

You can also perform operations, such as reboot, or force the system into an OFFLINE state and reload the default geolocation data files that the BIG-IP system uses to source the origin of a name resolution request.

About IP geolocation database updates

The BIG-IP® system uses an IP geolocation database to determine the origin of a name resolution request. The default database provides geolocation data for IPv4 addresses at the continent, country, state, ISP, and organization levels. The state-level data is worldwide and includes designations in other countries that correspond to the U.S. state-level in the geolocation hierarchy, for example, provinces in Canada. The default database also provides geolocation data for IPv6 addresses at the continent and country levels.

Note: You can access the ISP-level and organization-level geolocation data for IPv4 and IPv6 addresses using the iRules® `whereis` command.

Tip: If you require geolocation data at the city-level, contact your F5® Networks sales representative to purchase additional database files.

You can download a monthly update to the IP geolocation database from F5 Networks.

About Network Time Protocol (NTP)

Network Time Protocol (*NTP*) is a protocol that synchronizes the clocks on a network. Because DHCP is enabled for the BIG-IP® system by default, on the first boot, the BIG-IP system contacts your DHCP server and obtains the IP address of your NTP server. If the DHCP server provides this IP address, the NTP Device Configuration screen displays the NTP server information. If you do not have a DHCP server on your network, or if the DHCP server does not return the IP address of your NTP server, you can manually add the IP address of a NTP server to the BIG-IP system using the BIG-IP Configuration utility.

Configuring the NTP time server list

You can use the Configuration utility to specify a list of IP addresses of the servers that you want the BIG-IP® system to use when updating the time on network systems. You can also edit or delete the entries in the server list.

1. On the Main tab, click **System** > **Configuration** > **Device**.
    The NTP screen opens.
2. For the **Time Server List** setting, to add an IP address to the list:
    a) Type the IP address or host name of a time server in the **Address** field.
    b) Click **Add**.

3. For the **Time Server List** setting, to edit an IP address in the list:
    a) From the Time Server List, select an IP address.
        The IP address appears in the **Address** field.
    b) In the **Address** field, change the IP address.
    c) Click **Edit**.

4. For the **Time Server List** setting, to delete an IP address from the list:
    a) From the Time Server List, select an IP address.
        The IP address appears in the **Address** field.
    b) Click **Delete**.

5. Click **Update**.

About DNS configuration

Domain Name System (DNS) is an industry-standard, distributed Internet directory service that resolves domain names to IP addresses. When you enable DHCP, the system contacts your DHCP server to obtain the IP addresses of your local DNS servers and the domain names that the system searches to resolve local

host names. If the DHCP server provides this information, the DNS Device Configuration screen displays the information in the DNS Lookup Server List and the DNS Search Domain List.

If you do not have a DHCP server on your network, or if the DHCP server does not supply the information, you can manually create the two lists:

* The DNS Lookup Server List enables BIG-IP® system users to use IP addresses, host names, or fully-qualified domain names (FQDNs) for accessing virtual servers, nodes, or other network objects.
* The DNS Search Domain List enables BIG-IP systems to search for local domain lookups to resolve local host names.

Additionally, you can manually configure the *BIND Forwarder Server List* that provides DNS resolution for servers and other equipment load-balanced by the BIG-IP system (for the servers that the BIG-IP system uses for DNS proxy services).

Note: To use DNS Proxy services, you must enable the named service.

About local traffic properties

The BIG-IP® system includes a set of properties that apply globally to the local traffic management system. There are two categories of local traffic properties: General and Persistence. You can use the BIG-IP Configuration utility to configure and maintain these properties.

General local traffic properties

This table lists and describes global properties that you can configure to manage the behavior of the local traffic management system.

Property	Default value	Description
Auto Last Hop	Enabled (check box selected)	When selected (enabled), specifies that the system automatically maps the last hop for pools.
Maintenance Mode	Disabled (check box cleared)	When selected (enabled), specifies that the unit is in maintenance mode. In maintenance mode, the system stops accepting new connections and slowly completes the processing of existing connections.
VLAN-Keyed Connections	Enabled (check box selected)	Select this check box setting to enable VLAN-keyed connections. VLAN-keyed connections are used when traffic for the same connection must pass through the system several times, on multiple pairs of VLANs (or in different VLAN groups).
Path MTU Discovery	Enabled (check box selected)	When selected (enabled), specifies that the system discovers the maximum transmission unit (MTU) that it can send over a path without fragmenting TCP packets.
Reject Unmatched Packets	Enabled (check box selected)	Specifies that the BIG-IP® system sends a TCP RST packet in response to a non-SYN packet that matches a virtual server address and port or self IP address and port, but does not match an established connection. The BIG-IP system also sends a TCP RST packet in response to a packet matching a virtual server address or self IP address but specifying an invalid port. The TCP RST packet is sent on the client-side of the connection, and the source IP address of the reset is the relevant BIG-IP LTM® object

Property	Default value	Description
		address or self IP address for which the packet was destined. If you disable this setting, the system silently drops unmatched packets.
Eviction Policy	default-eviction-policy	Specifies the eviction policy for the system, which provides the system with guidelines for how aggressively it discards flows from the flow table. You can customize the eviction policy to prevent flow table attacks, where a large number of slow flows are used to negatively impact system resources. You can also set how the system responds to such flow problems in an eviction policy, and attach such eviction policies globally, to route domains, and to virtual servers, to protect the system, applications and network segments with a high level of customization.
SYN Check™ Activation Threshold	16384	Specifies the number of new or untrusted TCP connections that can be established before the system activates the SYN Cookies authentication method for subsequent TCP connections.
Layer 2 Cache Aging Time	300	Specifies, in seconds, the amount of time that records remain in the Layer 2 forwarding table, when the MAC address of the record is no longer detected on the network.
Share Single MAC Address	Disabled (check box cleared)	When this check box setting is cleared (disabled), the BIG-IP system assigns to each VLAN a unique MAC address that comes from a pool of available MAC addresses. If you create enough VLANs to exceed the number of MAC addresses available, the system then begins to assign the same MAC address to multiple VLANs. This is the default value and the most common configuration. When this check box setting is selected (enabled), the BIG-IP system causes all VLANs to share a single MAC address (global). This setting is equivalent to the BigDB variable `vlan.macassignment` and has two values, `unique` and `global`.
SNAT Packet Forwarding	TCP and UDP Only	Specifies the type of traffic for which the system attempts to forward (instead of reject) Any-IP packets, when the traffic originates from a member of a SNAT. There are two possible values: **TCP and UDP Only** specifies that the system forwards, for TCP and UDP traffic only, Any-IP packets originating from a SNAT member. **All Traffic** specifies that the system forwards, for all traffic types, Any-IP packets originating from a SNAT member.

Configuration Data Management

About BIG-IP system configuration data

When you perform configuration tasks on the BIG-IP® system, it generates underlying configuration data. The system stores this data so that the data is not lost when an unexpected system event occurs or when you restart the system. Before the system can store this data, however, the data must be saved.

The BIG-IP system configuration data exists in these states:

- The stored configuration comprises all of the configuration tasks that you have performed on the system and saved to the system configuration files.
- The running configuration comprises the stored configuration and all of the changes that you have made to the system since the last save operation. The BIG-IP system operates based on the running configuration.

Note: This information applies only to tmsh *users. When you use the Configuration utility, it manages and automatically saves all configuration data for you as you complete each configuration task; you do not need to perform any additional steps to save configuration data.*

About managing system configuration data using tmsh

When you use the Traffic Management Shell (tmsh) to configure the system, you must explicitly issue a save command to store the configuration data that you have generated. Otherwise, the newly-generated configuration data is not actually stored on the system. For more information about tmsh, see the *Traffic Management Shell (tmsh) Reference Guide*.

About the single configuration file (SCF)

A *single configuration file (SCF)* is a text file that contains the configuration of a BIG-IP® system. You can use this file to easily replicate the configuration across multiple BIG-IP systems. This not only saves you time, but also enables you to create a consistent, secure, comprehensive local traffic management environment on your network.

This sample shows some of the information contained in an SCF file:

```
vlan external {
    tag 4093
    interfaces 1.3
}
vlan internal {
    tag 4094
    interfaces 1.10
}
```

```
pool dev_https3 {
    members {
        10.60.10.105:https{}
        10.60.10.106:https{}
    }
}
```

Creating and saving an SCF

You can use `tmsh` to create and save a single configuration file (SCF).

Note: If you create an SCF file twice (on two different occasions), you can compare the contents of the two files.

1. Open the Traffic Management Shell (`tmsh`).
   `tmsh`
2. Create and save an SCF.
   `save sys config file [filename]`

   Note: If you include the `.scf` extension in the file name, the system does not add an additional file extension.

   The system gathers all of the commands that make up the running configuration, and then saves the configuration to a `.scf` file with the name you specify. By default, the system stores this file in the `/var/local/scf` directory, but you can specify a different path if you prefer.

Loading an SCF onto a target BIG-IP system

You can use `tmsh` to load a single configuration file (SCF) on one BIG-IP® system that you created on another BIG-IP system (hereafter referred to as the target BIG-IP system). This saves you from having to recreate the configuration multiple times. Loading an SCF resets the running configuration with the values contained in the stored configuration.

Important: If you run a `load` command or restart the system before you save your changes to the stored configuration, you will lose any changes.

Note: To successfully load a configuration that you have replicated, make sure that no line of the configuration is longer than 4096 characters. If there are more than 4096 characters in a single line, the system reverts to the previous running configuration.

1. Open the Traffic Management Shell (`tmsh`).
   `tmsh`
2. On the target BIG-IP system, load the saved SCF file.
   `tmsh load sys config file [filename]`
   The system saves the stored configuration to a backup file named `/var/local/scf/backup.scf`, and then uses the configuration stored in the SCF that you are loading.
3. Use a text editor to open the SCF and edit any data that is unique to the target BIG-IP system, such as the management IP address.
4. Save the SCF to the target BIG-IP system.
   `sys save config file [filename]`

If a backup SCF already exists, the system appends a number to the name of the existing backup file, and then creates a new backup file. In the case of multiple load and save operations:

- The first time the system backs up the running configuration during a load operation, the system names the backup file `/var/local/scf/backup.scf`.
- The next time the system backs up the running configuration, the system renames the file from `/var/local/scf/backup.scf` to `/var/local/scf/backup-1.scf` and creates a new file named `/var/local/scf/backup.scf`.
- If you run the `load` command a third time, the system renames the file from `/var/local/scf/backup-1.scf` to `/var/local/scf/backup-2.scf`, renames the `/var/local/scf/backup.scf` file to `/var/local/scf/backup-1.scf`, and again creates a new file named `/var/local/scf/backup.scf`.

Using an SCF to restore a BIG-IP system configuration

You can use `tmsh` to restore a BIG-IP® system configuration using either a specific single configuration file (SCF) or the factory default configuration.

1. Open the Traffic Management Shell (`tmsh`).
   `tmsh`

2. Restore the system configuration using one of these options:
   - Restore a system to the factory default configuration by using `tmsh load sys config default`. This command retains the management IP and the assigned root and administrator passwords. When you use this command, the system first saves the running configuration in the `backup.scf` file, and then resets the local traffic management and the operating system configuration to the factory default settings by loading the factory default SCF (`/defaults/defaults.scf`).
   - Restore a system with values defined in the specified SCF by using `tmsh load sys config file [filename]`. When you use this command, the system first saves the running configuration in the `backup.scf` file, and then resets the running configuration to the values contained in the specified SCF.

   Note: You must run the `save sys config partitions all` command to save the running configuration in the stored configuration files.

tmsh commands for single configuration files (SCFs)

You use `tmsh` to manage a single configuration file (SCF). This table lists an overview of `tmsh` commands used to manage SCF files.

tmsh command	Description
`save sys config file [filename]`	Saves a copy of the currently running configuration to an SCF. *Important: Saving a configuration to an SCF does not affect the running or stored configuration of the BIG-IP® system on which you run the command.*

tmsh command	Description
`load sys config file [filename]`	Replaces or restores an SCF with a saved configuration. When you use this command, the system saves any previously running configuration to the `/var/local/scf/` directory, by default.
`load sys config default`	Restores the factory default settings of the configuration file, while retaining the management IP address and the administrator user name and password.

Archives

About archives

When you initially configure the BIG-IP® system using the Setup utility and the BIG-IP Configuration utility, or `tmsh`, the system saves your configuration information. This information includes traffic management elements, such as virtual servers, pools, and profiles. Configuration data also consists of system and network definitions, such as interface properties, self IP addresses, VLANs, and more.

Once you have created the configuration data for the BIG-IP system, you can replicate all of this data in a separate file and then use this data later for these purposes:

Archive for disaster recovery
Using the Archives feature, you can back up the current configuration data, and if necessary, restore the data at a later time. F5® Networks recommends that you use this feature to mitigate the potential loss of BIG-IP system configuration data. To create an archive, you can use the BIG-IP Configuration utility, which stores the configuration data in a file known as a user configuration set, or UCS (`.ucs`) file. You can then use the UCS file to recover from any loss of data, in the unlikely event that you need to do so.

Propagate data to other systems
Using the single configuration file feature, you can quickly propagate the exact configuration of the BIG-IP system to other BIG-IP systems. To create a single configuration file, you export the configuration data to a file known as an SCF (`.scf`) file. You can then use the SCF file to configure another system in one simple operation.

By default, the system stores all archives in the `/var/local/ucs` directory. You can specify a different location, but if you do, the BIG-IP® Configuration utility does not display the UCS files when you view the archive list.

Before you replace a version of the BIG-IP system with a newer version, you should always create an *archive*, which is a backup copy of the configuration data. This archive is in the form of a user configuration set, or UCS. Then, if you need to recover that data later, you can restore the data from the archive that you created.

Important: To create, delete, upload, or download an archive, you must have either the Administrator or Resource Administrator role assigned to your user account.

About UCS files

A user configuration set, or UCS (`.ucs`) file, contains the following types of BIG-IP system configuration data:

- System-specific configuration files
- Product licenses
- User accounts and password information
- Domain name service (DNS) zone files

- Installed SSL keys and certificates

Each time you back up the configuration data, the BIG-IP system creates a new file with a `.ucs` extension. Each UCS file contains various configuration files needed for the BIG-IP system to operate correctly, as well as the configuration data.

About managing archives using the Configuration utility

When you create a new archive (or UCS file) using the Configuration utility, the BIG-IP® system automatically stores it at a default location, in the `/var/local/ucs` directory. You can create as many separate archives as you need, provided each archive has a unique file name. Also, you can specify that the BIG-IP system store an archive in a directory other than `/var/local/ucs`. In this case, however, the Configuration utility does not include the archive name in the list of archives on the Archives screen.

Creating and saving an archive using the Configuration utility

You can use the BIG-IP® Configuration utility to create and save archives on the BIG-IP system.

Important: Any UCS file that you create includes the host name of the BIG-IP system as part of the data stored in that file. Later, when you specify this UCS file while restoring configuration data to a BIG-IP system, the host name stored in this UCS file must match the host name of the system to which you are restoring the configuration data. Otherwise, the system does not fully restore the data. Also, if your configuration data includes SSL keys and certificates, make sure to store the archive file in a secure environment.

1. On the Main tab, click **System** > **Archives**.
   The Archives screen displays a list of existing UCS files.
2. Click **Create**.

   *Note: If the **Create** button is unavailable, you do not have permission to create an archive. You must have the Administrator role assigned to your user account.*

3. In the **File Name** field, type a unique file name for the archive.
   F5® recommends that the file name match the name of the BIG-IP system. For example, if the name of the BIG-IP system is `bigip2`, then the name of the archive file should be `bigip2.ucs`.
4. To encrypt the archive, for the **Encryption** setting, select **Enabled**.

   *Note: If the **Encryption** setting is unavailable, you must configure the **Archive Encryption** setting located on the Preferences screen.*

5. To include private keys, for the **Private Keys** setting, select **Include**.
   Make sure to store the archive file in a secure environment.
6. Click **Finished**.

Restoring data from an archive using the Configuration utility

In the unlikely event that the BIG-IP® system configuration data becomes corrupted, you can use the Configuration utility to restore data from an archive file. The /var/local/ucs directory is the only location on the BIG-IP system in which you can save and restore an archive. If no archive exists in that directory, then you cannot restore configuration data.

Important: *The host name stored in the archive file must match the host name of the BIG-IP system that you are restoring; otherwise, the system does not fully restore the data.*

---

1. On the Main tab, click **System** > **Archives**.
   The Archives screen displays a list of existing UCS files.

2. In the File Name column, click the name of the archive that you want to use to restore the configuration data.
   This displays the properties of that archive.

3. Click **Restore**.
   The system displays a progress message.

Viewing a list of existing archives using the Configuration utility

You can use the Configuration utility to view a list of archives that are stored in the the default directory, /var/local/ucs, on a BIG-IP® system. The Configuration utility displays the UCS file name, creation date, and file size.

On the Main tab, click **System** > **Archives**.
The Archives screen displays a list of existing UCS files.

Viewing archive properties using the Configuration utility

You can use the Configuration utility to view the properties of archives that are stored on the BIG-IP® system, including archive name, BIG-IP version, encryption state, creation date, and archive size.

1. On the Main tab, click **System** > **Archives**.
   The Archives screen displays a list of existing UCS files.

2. In the File Name column, click the name of the archive that you want to view.
   This displays the properties of that archive.

Downloading a copy of an archive to a management workstation

You can use the Configuration utility to download a copy of an archive to a management workstation. This provides an extra level of protection by preserving the configuration data on a remote system. In the unlikely event that you need to restore the data, and a BIG-IP® system event prevents you from accessing the archive in the BIG-IP system directory, you still have a backup copy of the configuration data.

1. On the Main tab, click **System** > **Archives**.

The Archives screen displays a list of existing UCS files.

2. In the File Name column, click the name of the archive that you want to view.

   This displays the properties of that archive.

3. For the **Archive File** setting, click the **Download: <filename>.ucs** button.

   A confirmation screen appears.

4. Click **Save**.

   The BIG-IP system downloads a copy of the UCS file to the system from which you initiated the download.

Uploading an archive from a management workstation

If you previously downloaded a copy of an archive to a management workstation, you can upload that archive to the BIG-IP® system at any time. This is useful when a BIG-IP system event has occurred that has caused the archive stored on the BIG-IP system to either become unavailable or corrupted.

You can use the Configuration utility to upload a copy of an archive stored on a management workstation.

Note: When you upload a copy of an archive, you must specify the exact path name for the directory in which the downloaded archive copy is stored.

1. On the Main tab, click **System** > **Archives**.
   The Archives screen displays a list of existing UCS files.

2. Click **Upload**.
   The Upload screen opens.

3. For the **File Name** setting, click **Browse**.

4. For the **Options** setting, select the **Overwrite existing archive file** check box if you want the BIG-IP system to overwrite any existing archive file.

   Note: The BIG-IP system overwrites an existing file with the uploaded file only when the name of the archive you are uploading matches the name of an archive on the BIG-IP system.

5. Click **Upload**.
   The specified archive is now uploaded to the /var/local/ucs directory on the BIG-IP system.

Deleting an archive using the Configuration utility

You can use the Configuration utility to delete an archive that is stored in the default UCS directory, /var/local/ucs, on the BIG-IP® system.

1. Open the Traffic Management Shell (tmsh).
   tmsh

2. Delete the specified archive file.
   delete sys ucs <filename>

   The specified UCS file is deleted.

About managing archives using tmsh

When you create a new archive using the Traffic Management Shell (`tmsh`), the BIG-IP® system automatically stores it at a default location, in the `/var/local/ucs` directory. You can create as many separate archives as you need, provided each archive has a unique file name. Also, you can specify that the BIG-IP system store an archive in a directory other than `/var/local/ucs`. In this case, however, `tmsh` does not include the archive name when you view a list of existing archives.

For more information about `tmsh` commands and options, see the man pages or the *Traffic Management Shell (tmsh) Reference Guide*.

Creating and saving an archive using tmsh

You can use `tmsh` to create and save archives (UCS files) on the BIG-IP® system.

Important: Any UCS file that you create includes the host name of the BIG-IP system as part of the data stored in that file. Later, when you specify this UCS file while restoring configuration data to a BIG-IP system, the host name stored in this UCS file must match the host name of the system to which you are restoring the configuration data. Otherwise, the system does not fully restore the data. Also, if your configuration data includes SSL keys and certificates, make sure to store the archive file in a secure environment.

1. Open the Traffic Management Shell (`tmsh`).
    `tmsh`
2. Save the running configuration of the system to a new UCS file, where <filename> is the name of the new UCS file.
    `save ucs <filename>`

Viewing a list of existing archives using tmsh

You can use `tmsh` to view a list of archives that are stored in the default directory, `/var/local/ucs`, on the BIG-IP® system.

1. Open the Traffic Management Shell (`tmsh`).
    `tmsh`
2. View a list of UCS files stored in `/var/local/ucs`.

    `show sys ucs`

    A list of UCS files displays.

Viewing archive properties using tmsh

You can use `tmsh` to view the properties of archives that are stored on the BIG-IP® system, including archive name, BIG-IP version, encryption state, creation date, and archive size.

1. Open the Traffic Management Shell (`tmsh`).

```
tmsh
```

2. View the properties for all UCS files stored in `/var/local/ucs`.

```
show sys ucs
```

Note: To view properties for a specific UCS file, include the UCS file name in the command sequence.

The properties for all UCS files displays.

Deleting an archive using tmsh

You can use `tmsh` to delete an archive that is stored in the default UCS directory, `/var/local/ucs`, on the BIG-IP® system.

1. Open the Traffic Management Shell (`tmsh`).
```
tmsh
```

2. Delete the specified UCS file.

```
delete sys ucs <filename>
```

The system deletes the specified UCS file.

About backing up and restoring archives using tmsh

After you have created an archive (UCS), you can use secure copy (SCP) to save a copy to a management workstation. This provides an extra level of protection by preserving the configuration data on a remote system. In the unlikely event that you need to restore the data and you are unable to access the archive in the BIG-IP® system directory, you still have a backup copy of the configuration data.

Important: If your configuration data includes SSL keys and certificates, make sure to store the archive file in a secure environment.

Once the UCS is in the `/var/local/ucs` directory, you can load and restore the archive data using `tmsh`.

Loading and restoring data from an archive using tmsh

In the unlikely event that the BIG-IP® system configuration data becomes corrupted, you can use `tmsh` to load and restore data from an archive file. The `/var/local/ucs` directory is the only location on the BIG-IP system from which you can restore an archive. If no archive exists in that directory, then you cannot restore configuration data.

Important: The host name stored in the archive file must match the host name of the BIG-IP system that you are restoring; otherwise, the system does not fully restore the data.

---

1. Open the Traffic Management Shell (`tmsh`).
```
tmsh
```

2. Load the configuration contained in a specifed UCS file, where <filename> is the name of the UCS file.

```
load ucs <filename>
```

The UCS is loaded into the running configuration of the system.

Disk Management

About disk management

You can manage the disk drives installed in the BIG-IP® system using the Configuration utility.

Adding a disk to an array

You can use the Configuration utility to add a disk to an array.

Note: Disk arrays apply only to BIG-IP® systems that contain hard disk drives (HDDs). These features are not available on systems that contain solid-state drives (SSDs).

1. On the Main tab, click **System** > **Disk Management**.
   The Disk Management screen displays a list of disks installed in the system.
2. For the **Software and Mixed Mode Disks** setting, select a Disk Name.
3. Click **Add Disk to Array**.
4. Confirm that you would like to add the selected disk to the array.

   Important: This destroys any data on the disk.

   When the disk is being added to the array, for the **Software and Mixed Mode Disks** area, the Array Status column changes to Replicating for that disk. When the replication process is complete, the Array Status column changes to OK.

Removing a disk from an array

You can use the Configuration utility to remove a disk from an array.

Warning: Removing a disk from the array destroys all data on the disk.

Note: Disk arrays apply only to BIG-IP® systems that contain hard disk drives (HDDs). These features are not available on systems that contain solid-state drives (SSDs).

1. On the Main tab, click **System** > **Disk Management**.
   The Disk Management screen displays a list of disks installed in the system.
2. For the **Software and Mixed Mode Disks** setting, select a Disk Name.
3. Click **Remove Disk From Array**.
4. Confirm that you would like to remove the selected disk from the array.

   Warning: This destroys any data on the disk.

When the disk is removed from the array, for the **Software and Mixed Mode Disks** area, the Array Status column changes to Undefined for that disk.

Software Management

About software management

You can manage the software images, hotfixes, and boot locations on the BIG-IP® system using the Configuration utility. You can also enable the automatic software update feature.

Importing a software image

If you previously downloaded a BIG-IP® software image file (ISO) to a management workstation, you can upload that file to the BIG-IP system.

You can use the Configuration utility to import an ISO that you have stored on a management workstation.

1. On the Main tab, click **System** > **Software Management** > **Image List**.
   The Image List screen displays a list of existing image files.
2. Click **Import**.
   The New Image screen opens.
3. For the **Software Image** setting, click **Browse**.
4. Click **Import**.
   A progress indicator displays as the BIG-IP system uploads the file.

   Note: Be sure that you do not navigate away from the screen until the image import process is complete.

Installing a software image

You can use the Configuration utility to install an ISO that you have imported to the BIG-IP® system.

1. On the Main tab, click **System** > **Software Management** > **Image List**.
   The Image List screen displays a list of existing image files.
2. For the **Available Images** setting, select the ISO to install.
   The Install Software Image screen opens.
3. For the **Select Disk** setting, select the disk on which to install the software (for example, MD1 or HD1).

   Note: You can install software only on inactive volumes. To install software to the active volume, you must boot to a different volume.

4. For the **Volume set name** setting, select the volume on which to install the software.
5. Click **Install**.
   A progress indicator displays as the BIG-IP system installs the software image.

Importing a software image

If you previously downloaded a BIG-IP® hotfix file to a management workstation, you can upload that file to the BIG-IP system.

You can use the Configuration utility to import hotfix that you have stored on a management workstation.

1. On the Main tab, click **System** > **Software Management** > **Hotfix List**.
   The Hotfix List screen displays a list of existing hotfix files.

2. Click **Import**.
   The Upload Hotfix screen opens.

3. For the **Software Image** setting, click **Browse**.

4. Click **Import**.
   A progress indicator displays as the BIG-IP system uploads the file.

   Note: Be sure that you do not navigate away from the screen until the image import process is complete.

Installing a hotfix

You can use the Configuration utility to install a hotfix that you have imported to the BIG-IP® system.

1. On the Main tab, click **System** > **Software Management** > **Hotfix List**.
   The Hotfix List screen displays a list of existing hotfix files.

2. For the **Available Images** setting, select the hotfix to install.
   The Install Software Hotfix screen opens.

3. For the **Select Disk** setting, select the disk on which to install the software (for example, MD1 or HD1).

   Note: You can install software only on inactive volumes. To install software to the active volume, you must boot to a different volume.

4. For the **Volume set name** setting, select the volume on which to install the software.

5. Click **Install**.
   A progress indicator displays as the BIG-IP system installs the hotfix.

Booting to a different volume

You can use the Configuration utility to boot to a different volume on the BIG-IP® system.

1. On the Main tab, click **System** > **Software Management** > **Boot Locations**.
   The Boot Locations List screen displays a list of available boot locations.

2. For the **Boot Location** setting, click a volume name (the target boot location).

3. For the **General Properties** setting for the target boot location, select whether to copy the configuration from the current boot location to the target boot location.

4. Click **Activate**.
   The system reboots to the selected volume.

About Liveinstall signature checking in ccmode

For each full release ISO, vADC OVA, and hotfix ISO, a corresponding signature file will be available with the .sig extension. The signature file is handled exactly like an ISO. When the ccmode feature is turned on, the installation process requires you to download the ISO file, as well as the `iso.sig`.

When you run the ccmode script to put the sensor into a Common Criteria configuration, a db variable called `liveinstall.checksig` is automatically enabled. This feature compares the ISO file against a sys software signature file, which is meant to catch integrity issues with the product.

Note: This feature can only be controlled through `tmsh`.

Signature validation is the first step performed during the liveinstall process, so if the corresponding signature file for the selected software is not in the library, the installation will not begin.

Important: If `liveinstall.checksig` is enabled, software installs will fail if the user copies only the ISO to the `/shared/images` directory. It is important to download both the ISO and the `iso.sig` files to the `/shared/images` directory.

In the event of liveinstall failure, two error messages can occur:

Signature file not found
This means you have not downloaded the corresponding `iso.sig` file with the ISO. The best way to verify if the `iso.sig` is present is to run the command `list sys software signature`. The command `show sys software` will not show the `iso.sig` files.

Archive signature test failed
This might happen if:

- The product ISO is in `/shared/images`.
- The `iso.sig` is present in `/shared/images`.
- When the `iso.sig` file was compared against the product ISO, the comparison failed.

In these instances, you will need to re-download the ISO and `iso.sig` files and try again.

Downloading the .sig file

If you are running your machine in ccmode, you will need to download an `iso.sig` file in addition to the ISO file that you normally download.

1. In a browser, open the F5® Downloads page (`https://downloads.f5.com`).
2. From the Downloads Overview page, click **Find a Download**.
    The Select a Product Line screen displays.
3. Download the version's base ISO file, such as version 11.5, and its associated signature file.
    The signature file has the same name as the ISO file with an additional .sig extension.
4. In `tmsh`, type: `modify sys db liveinstall.checksig value enable`
    You do not need a .sig file to install versions earlier than 11.5.
5. Type `install sys software image` and press Tab.
    Tab completion will only list the ISO files that have a corresponding `iso.sig` file present at the time the command was run.

Licensing

About licensing

You must activate a valid license on the BIG-IP® system before you can use it. To activate a license for the system, you must have a base registration key. The base registration key is a 27-character string that informs the license server about which F5® products you are entitled to license. The base registration key is preinstalled on your system. If you do not already have a base registration key, you can obtain one from F5 Technical Support (`http://support.f5.com/`).

If the system is not yet licensed, the Configuration utility prompts you to enter the base registration key. Certain systems might require you to enter keys for additional modules in the Add-On Registration Key List field.

Important: When you install a license for an Add-On module, you must also provision resources for the module.

Activating a license automatically using the Configuration utility

You can license the BIG-IP® system using the automatic activation method if the system is configured to route traffic to the Internet.

1. From a workstation attached to the network on which you configured the management interface, type the following URL syntax where `<management_IP_address>` is the address you configured for device management:
   `https://<management_IP_address>`
2. At the login prompt, type the default user name `admin`, and password `admin`, and click **Log in**.
   The Configuration utility opens. If this is the first time that you have run the Configuration utility, the system presents the Licensing screen of the Setup utility.
3. Click **Activate**.
4. For the **Activation Method** setting, select **Automatic**.
5. Click **Next**.

Activating a license manually using the Configuration utility

You can license a BIG-IP® system using the manual activation method if the system is not configured to route traffic to the Internet.

1. From a workstation attached to the network on which you configured the management interface, type the following URL syntax where `<management_IP_address>` is the address you configured for device management:
   `https://<management_IP_address>`
2. At the login prompt, type the default user name `admin`, and password `admin`, and click **Log in**.

The Configuration utility opens. If this is the first time that you have run the Configuration utility, the system presents the Licensing screen of the Setup utility.

3. Click **Activate**.

4. For the **Activation Method** setting, select **Manual**.

5. Click **Next**.

6. For the **Manual Method** setting, select **Copy/Paste Text**.

7. For the **Step 1: Dossier** setting, select and copy the dossier text.

8. Access the F5 license activation tool at `https://activate.f5.com/license/dossier.jsp`.

9. Paste the dossier text in the box, or if you saved the dossier as a text file, browse to that file.

10. Click **Next**.

The BIG-IP system reloads the configuration with the license. Traffic processing might be interrupted during this process.

Reactivating a license using the Configuration utility

You can reactivate an expired BIG-IP® system license.

1. From a workstation attached to the network on which you configured the management interface, type the following URL syntax where `<management_IP_address>` is the address you configured for device management:
   `https://<management_IP_address>`

2. At the login prompt, type the default user name `admin`, and password `admin`, and click **Log in**.
   The Configuration utility opens. If this is the first time that you have run the Configuration utility, the system presents the Licensing screen of the Setup utility.

3. On the Main tab, click **System** > **License**.
   The License Summary screen displays.

4. Click **Re-activate**.

5. For the **Activation Method** setting, select either **Automatic** or **Manual**.

6. Click **Next**.

Resource Provisioning

About resource provisioning

You can manage the provisioning of system memory, disk space, and CPU usage among licensed modules on the BIG-IP® system.

There are four available resource allocation settings for modules.

None/Disabled
Specifies that a module is not provisioned. A module that is not provisioned does not run.

Dedicated
Specifies that the system allocates all CPU, memory, and disk resources to one module. When you select this option, the system sets all other modules to None (Disabled).

Nominal
Specifies that, when first enabled, a module gets the least amount of resources required. Then, after all modules are enabled, the module gets additional resources from the portion of remaining resources.

Minimum
Specifies that when the module is enabled, it gets the least amount of resources required. No additional resources are ever allocated to the module.

Provisioning the BIG-IP system using the Configuration utility

After you have activated a license on the BIG-IP® system, you can use the Configuration utility to provision the licensed modules.

1. On the Main tab, click **System** > **Resource Provisioning**.
2. For licensed modules, select either Minimum or Nominal, as needed.
3. Click **Submit**.
4. Reboot the system:
   a) On the Main tab, click **System** > **Configuration** > **Device** > **General**.
   b) Click **Reboot**.

Provisioning the BIG-IP system using tmsh

You can use tmsh to provision the BIG-IP® system. For more information about the tmsh provision options, type help sys provision.

Important: You must provision the BIG-IP system before you configure it; otherwise, you lose the system configuration when you provision the system.

1. Open the Traffic Management Shell (`tmsh`).

   ```
   tmsh
   ```

2. View the current provisioning of the system.

   ```
   list sys provision
   ```

   The system displays the provision configuration. In this example, the system has nominal provisioning for LTM® and the other modules are unprovisioned.

   ```
   sys provision avr { }
   sys provision gtm { }
   sys provision lc { }
   sys provision ltm {
       level nominal
   }
   ```

3. Modify the current provisioning of the system using these parameters.

   ```
   modify sys provision <module_name> level <level_type>
   ```

   This example sets nominal provisioning for AVR. `modify sys provision avr level nominal` The system displays the provision configuration. In this example, the system now has nominal provisioning for LTM and AVR.

   ```
   sys provision avr { }
   sys provision gtm { }
   sys provision lc { }
   sys provision ltm {
       level nominal
   }
   ```

4. Save the changes to the stored configuration.

   ```
   save sys config
   ```

5. Verify the current provisioning of the system.

   ```
   list sys provision
   ```

About FPGA firmware selection

Some BIG-IP® systems support configuring the type of FPGA firmware that is enabled. You can choose from these types, depending on the hardware platform and provisioned modules.

Important: Changing the firmware type when using a multi-blade VIPRION® C2400 chassis is not recommended, as you might not get the expected scalable performance.

standard-balanced-fpga
   Suitable for all uses.

l4-performance-fpga
   Suitable for CGNAT standalone and BIG-IP Local Traffic Manager™ (LTM®) L4-centric deployments.

l7-intelligent-fpga
   Suitable for BIG-IP DNS request hardware DDoS validation and DNS hardware response cache.

Selecting an FPGA firmware type

You can use the Configuration utility to select a specific FPGA firmware type.

1. On the Main tab, click **System** > **Resource Provisioning**.
2. For the **FPGA Firmware Selection** setting, select the check box for the FPGA firmware type that you would like to use.
3. Click **Submit**.

Supported platforms for FPGA firmware selection

These platforms support FPGA firmware selection.

Platform family	Platform model
VIPRION®	B2250 blade
VIPRION	C2200 chassis
VIPRION	C2400 chassis

Platform Properties

About platform properties

Part of managing a BIG-IP® system involves configuring and maintaining a certain set of system properties. These properties consist of general platform properties, such as the BIG-IP system host name, IP address, and passwords for its system administrative accounts.

About general properties

You can configure these general properties for the BIG-IP® system platform:

The management port and TMM
The BIG-IP system has a management port to handle administrative traffic, and TMM switch interfaces to handle application traffic. *TMM switch interfaces* are those interfaces controlled by the Traffic Management Microkernel (TMM) service.

Management port configuration
By default, DHCP is disabled for the management port on the BIG-IP system. When enabled, DHCP uses UDP ports 67 and 68. On the first boot, the BIG-IP system contacts your DHCP server and obtains a lease for an IP address and default route for the management port, and DNS and NTP servers. You must then configure other system attributes, such as host name and domain name servers. When DHCP is disabled, you manually configure the management port by assigning an IP address and netmask to the port. The IP address that you assign to the management port must be on a different network than the self IP addresses that you assign to VLANs. You can use either an IPv4 or an IPv6 address for the management port. Additionally, if you intend to manage the BIG-IP system from a node on a different subnet of your network, you can specify an IP address for the BIG-IP system to use as a default route to the management port.

Note: If you do not have a DHCP server on your network, the BIG-IP system assigns a default IP address of `192.168.1.245` to the management port of appliances and virtual systems, and `192.186.1.246` to the management port of VIPRION® systems.

Host name
Every BIG-IP system must have a host name that is a fully qualified domain name (FQDN). An example of a host name is `bigip-02.win.net`.

Host IP address
Every BIG-IP system must have a host IP address. This IP address can be the same as the address that you used for the management port, or you can assign a unique address. The default value on the screen for this setting is **Use Management Port IP Address**.

Time zone
Another of the general platform properties that you can specify is the time zone. The many time zones that you can choose from are grouped into these categories: Africa, America, Antarctica, Arctic, Asia,

Atlantic, Australia, Europe, Indian, and Pacific. You should specify the time zone region that most closely represents the location of the BIG-IP system you are configuring.

About redundant device properties

A BIG-IP® system is typically part of a device group that synchronizes configuration data across two or more BIG-IP devices and provides high availability (failover and connection mirroring).

To ensure that this operates successfully, you assign a device group (to the `root` folder) to which you want to synchronize configuration data. All folders and sub-folders in the folder hierarchy inherit this device group as a folder attribute.

You also assign a floating traffic group to the `root` folder. All folders and sub-folders in the folder hierarchy inherit this traffic group as a folder attribute.

About user administration properties

Part of managing platform-related properties is maintaining passwords for the system account. You can also configure the system to allow certain IP addresses to access the BIG-IP® system through SSH.

About administrative account passwords

When you ran the Setup utility on the BIG-IP® system, you set up some administrative accounts. Specifically, you set up the `root` and `admin` accounts. The `root` and `admin` accounts are for use by BIG-IP system administrators.

Users logging in with the `root` account have terminal and browser access to the BIG-IP system. By default, users logging in with the `admin` account have browser-only access to the BIG-IP system. You can use the general screen for platform properties to change the passwords for `root` and `admin` accounts on a regular basis. To change a password, locate the **Root Account** or **Admin Account** setting, and in the **Password** field, type a new password. In the **Confirm** field, re-type the same password.

About SSH access configuration

When you configure SSH access, you enable user access to the BIG-IP® system through SSH. Also, only the IP addresses that you specify are allowed access to the system using SSH.

Configuring platform properties

You can use the Configuration utility to configure the platform properties of the BIG-IP® system.

1. On the Main tab, click **System** > **Platform**.
   The Platform screen opens.

2. Configure the management port:
   a) In the General Properties area, for the **Management Port Configuration** setting, select either **Automatic (DHCP)** or **Manual**.

3. In the Redundant Device Properties area, for the **Root Folder Traffic Group** setting, select a device group to which you want to synchronize configuration data.

4. Configure the root and admin account passwords:
   a) In the User Administration area, for the **Root Account** setting, type a new password in the **Password** field and re-type the new password in the **Confirm** field.
   b) For the **Admin Account** setting, type a new password in the **Password** field and re-type the new password in the **Confirm** field.

5. Configure SSH access to the BIG-IP system:
   a) In the User Administration area, select the **Enabled** check box for the **SSH Access** setting.
   b) For the **SSH IP Allow** setting, select either *** All Addresses** or **Specify Range**, which enables you to specify a range of addresses for which access is allowed.

6. Click **Update**.

High Availability Fail-safe

About system fail-safe

When you configure system fail-safe, the BIG-IP® system monitors various hardware components, as well as the heartbeat of various system services, and can take action if the system detects a heartbeat failure.

You can configure the BIG-IP system to monitor the switch board component and then take some action if the BIG-IP system detects a failure. Using the BIG-IP Configuration utility, you can specify the action that you want the BIG-IP system to take when the component fails. The BIG-IP system can perform these actions:

- Reboot the BIG-IP system.
- Restart all system services.
- Go offline.
- Go offline and cancel the TMM service.
- Fail over and restart TMM.

Configuring system fail-safe

You can use the BIG-IP® Configuration utility to configure system fail-safe for the BIG-IP system. You configure system fail-safe when you want the system to take a specific action when it detects either a switch board failure or a heartbeat failure on a particular system service.

1. On the Main tab, click **System** > **High Availability** > **Fail-safe** > **System**.
   The System screen opens.
2. In the System Trigger Properties area, for the **Switch Board Failure** setting, select the action that you want the BIG-IP system to take in the event of a switch board component failure.
3. If you want to configure a heartbeat failure action for one or more system services, then in the System Services area:
   a) In the Name column, click a service name.
   b) From the **Heartbeat** list, select **Enabled**.
   c) From the **Heartbeat Failure Action** list, select the action that you want the system to take when the system detects a heartbeat failure for the service.
   d) Repeat these steps for each service for which you want to configure a heartbeat failure action.

4. Click **Update**.

About gateway fail-safe

One type of network failure detection is known as gateway fail-safe, which applies to redundant system configurations only. *Gateway fail-safe* monitors traffic between an active BIG-IP® system in a device group and a pool containing a gateway router. You configure the gateway fail-safe feature if you want the BIG-IP

system to take an action, such as fail over, when some number of gateway routers in a pool of routers becomes unreachable.

You can configure gateway fail-safe using the BIG-IP Configuration utility. Configuring gateway fail-safe means designating a pool of routers as a gateway fail-safe pool. When you designate a pool as a gateway fail-safe pool, you provide this information:

- The name of the pool.
- The name of a BIG-IP device in a device group (either the local device or any other device group member).
- The minimum number of gateway pool members that must be available to avoid the designated action.
- The action that the BIG-IP system should take when the number of available gateway pool members drops below the designated threshold. The default value is **Failover**.

After you configure gateway fail-safe, by specifying an action of **Failover**, the named BIG-IP device (and only that device) fails over to another device group member whenever the number of available pool members falls below the specified threshold. Although all device group members share their pool configurations, each device ignores any gateway fail-safe configuration that does not specify itself as the device associated with the specified gateway pool.

Configuring gateway fail-safe

Before you can configure gateway fail-safe, you must have already created a gateway pool for the BIG-IP® system to use to forward traffic.

When you configure gateway fail-safe, the system monitors traffic between a specified BIG-IP device group member and the specified gateway pool. If the number of available pool members falls below the specified threshold, the system takes the specified action. You can use the BIG-IP Configuration utility to configure gateway fail-safe for the BIG-IP system.

1. On the Main tab, click **System** > **High Availability** > **Fail-safe** > **Gateway**.
   The Gateway screen opens.
2. Click **Add**.
3. In the Configuration area, for the **Gateway Pool** setting, select a pool to use as the gateway fail-safe pool.
4. For the **Device** setting, select a device name.
5. For the **Threshold** setting, specify the minimum number of gateway pool members that must be available.
   The system triggers the gateway fail-safe action if the threshold falls below this value.
6. For the **Action** setting, select the action that the system takes if the threshold falls below the minimum available members.
7. Click **Finished**.

About VLAN fail-safe

For maximum reliability, the BIG-IP® system supports failure detection on all VLANs. When you configure the fail-safe option for a VLAN, the BIG-IP system monitors network traffic going through that VLAN. If the BIG-IP system detects a loss of traffic on the VLAN and the fail-safe timeout period has elapsed, the BIG-IP system attempts to generate traffic by issuing ARP requests to nodes accessible through the VLAN. The BIG-IP system also generates an ARP request for the default route, if the default router is accessible

from the VLAN. Failover is averted if the BIG-IP system is able to send and receive any traffic on the VLAN, including a response to its ARP request.

For a redundant system configuration, if the BIG-IP system does not receive traffic on the VLAN before the timeout period expires, the system can initiate failover to another device group member, reboot, or restart all system services. For a single device configuration, the system can either reboot or restart all system services. The default action for both configurations is **Reboot**.

Warning: *You should configure the fail-safe option on a VLAN only after the BIG-IP system is in a stable production environment. Otherwise, routine network changes might cause failover unnecessarily.*

Each interface card installed on the BIG-IP system is typically mapped to one or more different VLANs. Thus, when you set the fail-safe option on a particular VLAN, you need to know the interface to which the VLAN is mapped. You can use the BIG-IP Configuration utility to view VLAN names and their associated interfaces.

Configuring VLAN fail-safe

Before you can configure VLAN fail-safe, you must have already created a VLAN for the BIG-IP® system.

You configure fail-safe for a specific VLAN when you want the system to monitor traffic for that VLAN. If the system detects no traffic on the VLAN after some number of seconds has elapsed, the system takes the specified action. You can use the BIG-IP Configuration utility to configure VLAN fail-safe.

1. On the Main tab, click **System** > **High Availability** > **Fail-safe** > **VLANs**.
   The VLANs screen opens.
2. Click **Add**.
3. In the Configuration area, for the **VLAN** setting, select a VLAN.
4. For the **Timeout** setting, specify the number of seconds that a system can run without detecting network traffic on this VLAN before it takes the fail-safe action.
5. For the **Action** setting, select the action that the system takes when it does not detect any traffic on this VLAN.
6. Click **Finished**.

Compression

About data compression strategies

If you need more control over the way that the BIG-IP® system compresses data than what the standard HTTP compression profile configuration provides, you can use the Traffic Management Shell (tmsh) to enable a compression strategy other than the default strategy.

You can choose from five compression strategies: Latency, Speed, Size, Ratio, and Adaptive. The default compression strategy is Latency. The BIG-IP system uses the compression strategy that you select to determine which compression provider (hardware compression or software compression) to use for a given HTTP response. Once an HTTP response is assigned to a compression provider, the response remains associated with that compression provider until the response completes.

Setting the data compression strategy

You can use tmsh to set a compression strategy other than the default. The default compression strategy is Latency.

1. Open the Traffic Management Shell (tmsh).

   tmsh

2. Set a specified compression strategy.

   ```
   sys modify db compression.strategy value [ latency | speed | size | ratio |
   adaptive ]
   ```

   This example enables the Adaptive compression strategy: sys modify db compression.strategy value adaptive

Available compression strategies

When using tmsh to configure compression for a BIG-IP® system, you can choose from these compression strategies.

Compression strategy	Description
Latency	This is the default compression strategy. The system favors the latency of compression providers and delays selection of a provider until data arrives. This strategy helps to better distribute the workload placed on each provider. Since each provider has different compression capabilities (for example, different potential throughput), this strategy focuses on the total throughput of the device by keeping metrics on the current throughput level of all requests combined. As the device approaches its theoretical work limit, the provider becomes less favorable by the strategy, and new work is assigned to the least busy provider.

Compression strategy	Description
	The Latency strategy provides a better throttle mechanism when a large number of compression requests are in the queue.
Speed	The system uses the hardware compression provider to the fullest extent possible. Only when the hardware is busy does the system use a software compression provider to compress HTTP server responses. The Speed strategy is best used for bulk compression and for limiting CPU overhead.
Size	The system performs as much compression in the software as possible using a ratio of TMM and Offload. Only when the software is busy does the system use the hardware compression provider to compress HTTP server responses. The Size strategy gives the best ratio at the expense of CPU overhead.
Ratio	The system uses a weighted round robin approach to decide which compression provider to use to compress data. The Ratio strategy limits CPU overhead while giving good compression ratios.
Adaptive	The system first uses a software compression provider to compress HTTP server responses. The system switches to the hardware compression providers based on both the gzip compression level that you set in the HTTP compression profile and the hardware compression provider that the system contains. As load on the system increases, the system responds by reducing the desired gzip compression level (specified in the HTTP compression profile). The system uses the hardware compression provider only when that provider can deliver the specified or systematically-reduced gzip compression level. The Adaptive strategy gives you the most control over how LTM® handles compression.

Legal Notices and Acknowledgments

Legal Notices

Publication Date

This document was published on November 28, 2017.

Publication Number

MAN-0529-00

Copyright

Trademarks

Patents

This product may be protected by one or more patents indicated at:
http://www.f5.com/about/guidelines-policies/patents

Link Controller Availability

This product is not currently available in the United States.

Export Regulation Notice

This product may include cryptographic software. Under the Export Administration Act, the United States government may consider it a criminal offense to export this product from the United States.

RF Interference Warning

This is a Class A product. In a domestic environment this product may cause radio interference, in which case the user may be required to take adequate measures.

FCC Compliance

This equipment has been tested and found to comply with the limits for a Class A digital device pursuant to Part 15 of FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This unit generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user, at his own expense, will be required to take whatever measures may be required to correct the interference.

Any modifications to this device, unless expressly approved by the manufacturer, can void the user's authority to operate this equipment under part 15 of the FCC rules.

Canadian Regulatory Compliance

This Class A digital apparatus complies with Canadian ICES-003.

Standards Compliance

This product conforms to the IEC, European Union, ANSI/UL and Canadian CSA standards applicable to Information Technology products at the time of manufacture.

Acknowledgments

This product includes software developed by Gabriel Forté.

This product includes software developed by Bill Paul.

This product includes software developed by Jonathan Stone.

This product includes software developed by Manuel Bouyer.

This product includes software developed by Paul Richards.

This product includes software developed by the NetBSD Foundation, Inc. and its contributors.

This product includes software developed by the Politecnico di Torino, and its contributors.

This product includes software developed by the Swedish Institute of Computer Science and its contributors.

This product includes software developed by the University of California, Berkeley and its contributors.

This product includes software developed by the Computer Systems Engineering Group at the Lawrence Berkeley Laboratory.

This product includes software developed by Christopher G. Demetriou for the NetBSD Project.

This product includes software developed by Adam Glass.

This product includes software developed by Christian E. Hopps.

This product includes software developed by Dean Huxley.

This product includes software developed by John Kohl.

This product includes software developed by Paul Kranenburg.

This product includes software developed by Terrence R. Lambert.

This product includes software developed by Philip A. Nelson.

This product includes software developed by Herb Peyerl.

This product includes software developed by Jochen Pohl for the NetBSD Project.

This product includes software developed by Chris Provenzano.

This product includes software developed by Theo de Raadt.

This product includes software developed by David Muir Sharnoff.

This product includes software developed by SigmaSoft, Th. Lockert.

This product includes software developed for the NetBSD Project by Jason R. Thorpe.

This product includes software developed by Jason R. Thorpe for And Communications, http://www.and.com.

This product includes software developed for the NetBSD Project by Frank Van der Linden.

This product includes software developed for the NetBSD Project by John M. Vinopal.

This product includes software developed by Christos Zoulas.

This product includes software developed by the University of Vermont and State Agricultural College and Garrett A. Wollman.

This product includes software developed by Balazs Scheidler (bazsi@balabit.hu), which is protected under the GNU Public License.

This product includes software developed by Niels Mueller (nisse@lysator.liu.se), which is protected under the GNU Public License.

In the following statement, "This software" refers to the Mitsumi CD-ROM driver: This software was developed by Holger Veit and Brian Moore for use with 386BSD and similar operating systems. "Similar operating systems" includes mainly non-profit oriented systems for research and education, including but not restricted to NetBSD, FreeBSD, Mach (by CMU).

This product includes software developed by the Apache Group for use in the Apache HTTP server project (http://www.apache.org/).

This product includes software licensed from Richard H. Porter under the GNU Library General Public License (© 1998, Red Hat Software), www.gnu.org/copyleft/lgpl.html.

This product includes the standard version of Perl software licensed under the Perl Artistic License (© 1997, 1998 Tom Christiansen and Nathan Torkington). All rights reserved. You may find the most current standard version of Perl at http://www.perl.com.

This product includes software developed by Jared Minch.

This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (http://www.openssl.org/).

This product includes cryptographic software written by Eric Young (eay@cryptsoft.com).

This product contains software based on oprofile, which is protected under the GNU Public License.

This product includes software with glib library utility functions, which is protected under the GNU Public License.

This product includes software with grub2 bootloader functions, which is protected under the GNU Public License.

This product includes software with the Intel Gigabit Linux driver, which is protected under the GNU Public License. Copyright ©1999 - 2012 Intel Corporation.

This product includes software with the Intel 10 Gigabit PCI Express Linux driver, which is protected under the GNU Public License. Copyright ©1999 - 2012 Intel Corporation.

This product includes RRDtool software developed by Tobi Oetiker (http://www.rrdtool.com/index.html) and licensed under the GNU General Public License.

This product contains software licensed from Dr. Brian Gladman under the GNU General Public License (GPL).

This product includes software developed by the Apache Software Foundation (http://www.apache.org/).

This product includes Hypersonic SQL.

This product contains software developed by the Regents of the University of California, Sun Microsystems, Inc., Scriptics Corporation, and others.

This product includes software developed by the Internet Software Consortium.

This product includes software developed by Nominum, Inc. (http://www.nominum.com).

This product contains software developed by Broadcom Corporation, which is protected under the GNU Public License.

This product contains software developed by MaxMind LLC, and is protected under the GNU Lesser General Public License, as published by the Free Software Foundation.

This product includes software under license from Qosmos (www.qosmos.com).

This product includes software developed by Andrew Tridgell, which is protected under the GNU Public License, copyright ©1992-2000.

This product includes software developed by Jeremy Allison, which is protected under the GNU Public License, copyright ©1998.

This product includes software developed by Guenther Deschner, which is protected under the GNU Public License, copyright ©2008.

This product includes software developed by www.samba.org, which is protected under the GNU Public License, copyright ©2007.

This product includes software from Allan Jardine, distributed under the MIT License.

This product includes software from Trent Richardson, distributed under the MIT License.

This product includes vmbus drivers distributed by Microsoft Corporation.

This product includes software from Cavium.

This product includes software from Webroot, Inc.

This product includes software from Maxmind, Inc.

This product includes software from OpenVision Technologies, Inc. Copyright ©1993-1996, OpenVision Technologies, Inc. All Rights Reserved.

This product includes software developed by Matt Johnson, distributed under the MIT License. Copyright ©2012.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

This product includes Intel QuickAssist kernel module, library, and headers software licensed under the GNU General Public License (GPL).

This product includes gd-libgd library software developed by the following in accordance with the following copyrights:

- Portions copyright ©1994, 1995, 1996, 1997, 1998, 2000, 2001, 2002 by Cold Spring Harbor Laboratory. Funded under Grant P41-RR02188 by the National Institutes of Health.
- Portions copyright ©1996, 1997, 1998, 1999, 2000, 2001, 2002 by Boutell.Com, Inc.
- Portions relating to GD2 format copyright ©1999, 2000, 2001, 2002 Philip Warner.
- Portions relating to PNG copyright ©1999, 2000, 2001, 2002 Greg Roelofs.
- Portions relating to gdttf.c copyright ©1999, 2000, 2001, 2002 John Ellson (ellson@lucent.com).
- Portions relating to gdft.c copyright ©2001, 2002 John Ellson (ellson@lucent.com).
- Portions copyright ©2000, 2001, 2002, 2003, 2004, 2005, 2006, 2007 2008 Pierre-Alain Joye (pierre@libgd.org).
- Portions relating to JPEG and to color quantization copyright ©2000, 2001, 2002, Doug Becker and copyright ©1994, 1995, 1996, 1997, 1998, 1999, 2000, 2001, 2002, Thomas G. Lane. This software is based in part on the work of the Independent JPEG Group.
- Portions relating to WBMP copyright 2000, 2001, 2002 Maurice Szmurlo and Johan Van den Brande. Permission has been granted to copy, distribute and modify gd in any context without fee, including a commercial application, provided that this notice is present in user-accessible supporting documentation.

This product includes software developed by Oracle America, Inc. Copyright ©2012.

1. Java Technology Restrictions. Licensee shall not create, modify, change the behavior of, or authorize licensees of licensee to create, modify, or change the behavior of, classes, interfaces, or subpackages that are in any way identified as "java", "javax", "sun" or similar convention as specified by Oracle in any naming convention designation. In the event that Licensee creates an additional API(s) which: (a) extends the functionality of a Java Environment; and (b) is exposed to third party software developers for the purpose of developing additional software which invokes such additional API, Licensee must promptly publish broadly an accurate specification for such API for free use by all developer.
2. Trademarks and Logos. This License does not authorize an end user licensee to use any Oracle America, Inc. name, trademark, service mark, logo or icon. The end user licensee acknowledges that Oracle owns the Java trademark and all Java-related trademarks, logos and icon including the Coffee Cup and Duke ("Java Marks") and agrees to: (a) comply with the Java Trademark Guidelines at http://www.oraclc.com/html/3party.html; (b) not do anything harmful to or inconsistent with Oracle's rights in the Java Marks; and (c) assist Oracle in protecting those rights, including assigning to Oracle any rights acquired by Licensee in any Java Mark.
3. Source Code. Software may contain source code that, unless expressly licensed for other purposes, is provided solely for reference purposes pursuant to the terms of your license. Source code may not be redistributed unless expressly provided for in the terms of your license.
4. Third Party Code. Additional copyright notices and license terms applicable to portion of the Software are set forth in the THIRDPARTYLICENSEREADME.txt file.
5. Commercial Features. Use of the Commercial Features for any commercial or production purpose requires a separate license from Oracle. "Commercial Features" means those features identified in Table I-I (Commercial Features In Java SE Product Editions) of tile Software documentation accessible at http://www.oracle.com/technetwork/java/javase/documentation/index.html.

This product includes utilities developed by Linus Torvalds for inspecting devices connected to a USB bus.

This product includes perl-PHP-Serialization software, developed by Jesse Brown, copyright ©2003, and distributed under the Perl Development Artistic License (http://dev.perl.org/licenses/artistic.html).

This product includes software developed by members of the CentOS Project under the GNU Public License, copyright ©2004-2011 by the CentOS Project.

This product includes software developed by members of the OpenJDK Project under the GNU Public License Version 2, copyright ©2012 by Oracle Corporation.

This product includes software developed by The VMWare Guest Components Team under the GNU Public License Version 2, copyright ©1999-2011 by VMWare, Inc.

This product includes software developed by The Netty Project under the Apache Public License Version 2, copyright ©2008-2012 by The Netty Project.

This product includes software developed by Stephen Colebourne under the Apache Public License Version 2, copyright ©2001-2011 Joda.org.

This product includes software developed by the GlassFish Community under the GNU Public License Version 2 with classpath exception, copyright ©2012 Oracle Corporation.

This product includes software developed by the Mort Bay Consulting under the Apache Public License Version 2, copyright ©1995-2012 Mort Bay Consulting.

This product contains software developed by members of the Jackson Project under the GNU Lesser General Public License Version 2.1, ©2007 – 2012 by the Jackson Project".

This product contains software developed by QOS.ch under the MIT License, ©2004 – 2011 by QOS.ch.

This product includes software licensed from Gerald Combs (gerald@wireshark.org) under the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or any later version. Copyright ©1998 Gerald Combs.

This product includes software developed by Daniel Stenberg. Copyright ©1996 - 2012, Daniel Stenberg, (daniel@haxx.se). All rights reserved.

Permission to use, copy, modify, and distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF THIRD PARTY RIGHTS. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Except as contained in this notice, the name of a copyright holder shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Software without prior written authorization of the copyright holder.

This product includes software licensed from Rémi Denis-Courmont under the GNU Library General Public License. Copyright ©2006 - 2011.

This product includes software developed by jQuery Foundation and other contributors, distributed under the MIT License. Copyright ©2014 jQuery Foundation and other contributors (http://jquery.com/).

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Permission to use, copy, and distribute this software and its documentation for any purpose with or without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation. Permission to modify the software is granted, but not the right to distribute the complete modified source code. Modifications are to be distributed as patches to the released version. Permission to distribute binaries produced by compiling modified sources is granted, provided you

1. distribute the corresponding source modifications from the released version in the form of a patch file along with the binaries,
2. add special version identification to distinguish your version in addition to the base release version number,
3. provide your name and address as the primary contact for the support of your modified version, and
4. retain our contact information in regard to use of the base software.

Permission to distribute the released version of the source code along with corresponding source modifications in the form of a patch file is granted with same provisions 2 through 4 for binary distributions. This software is provided "as is" without express or implied warranty to the extent permitted by applicable law.

This product includes software developed by Brian Gladman, Worcester, UK Copyright ©1998-2010. All rights reserved. The redistribution and use of this software (with or without changes) is allowed without the payment of fees or royalties provided that:

- source code distributions include the above copyright notice, this list of conditions and the following disclaimer;
- binary distributions include the above copyright notice, this list of conditions and the following disclaimer in their documentation.

This software is provided "as is" with no explicit or implied warranties in respect of its operation, including, but not limited to, correctness and fitness for purpose.

This product includes software developed by the Computer Systems Engineering Group at Lawrence Berkeley Laboratory. Copyright ©1990-1994 Regents of the University of California. All rights reserved. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgment: This product includes software developed by the Computer Systems Engineering Group at Lawrence Berkeley Laboratory.
4. Neither the name of the University nor of the Laboratory may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This product includes software developed by Sony Computer Science Laboratories Inc. Copyright © 1997-2003 Sony Computer Science Laboratories Inc. All rights reserved. Redistribution and use in source

and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY SONY CSL AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL SONY CSL OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This product contains software developed by Google, Inc. Copyright ©2011 Google, Inc.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

This software incorporates JFreeChart, ©2000-2007 by Object Refinery Limited and Contributors, which is protected under the GNU Lesser General Public License (LGPL).

This product contains software developed by the Mojarra project. Source code for the Mojarra software may be obtained at https://javaserverfaces.dev.java.net/.

This product includes software developed by McAfee®.

This product includes software developed by Ian Gulliver ©2006, which is protected under the GNU General Public License, as published by the Free Software Foundation.

This product contains software developed by the RE2 Authors. Copyright ©2009 The RE2 Authors. All rights reserved. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

• Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
• Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
• Neither the name of Google Inc. nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE

This product includes the Zend Engine, freely available at http://www.zend.com.

This product includes software developed by Digital Envoy, Inc.

This product contains software developed by NuSphere Corporation, which is protected under the GNU Lesser General Public License.

This product contains software developed by Erik Arvidsson and Emil A Eklund.

This product contains software developed by Aditus Consulting.

This product contains software developed by Dynarch.com, which is protected under the GNU Lesser General Public License, version 2.1 or later.

This product contains software developed by InfoSoft Global (P) Limited.

This product includes software written by Steffen Beyer and licensed under the Perl Artistic License and the GPL.

This product includes software written by Makamaka Hannyaharamitu ©2007-2008.

Rsync was written by Andrew Tridgell and Paul Mackerras, and is available under the GNU Public License.

This product includes Malloc library software developed by Mark Moraes. (©1988, 1989, 1993, University of Toronto).

This product includes open SSH software developed by Tatu Ylonen (ylo@cs.hut.fi), Espoo, Finland (©1995).

This product includes open SSH software developed by Niels Provos (©1999).

This product includes SSH software developed by Mindbright Technology AB, Stockholm, Sweden, www.mindbright.se, info@mindbright.se (©1998-1999).

This product includes free SSL software developed by Object Oriented Concepts, Inc., St. John's, NF, Canada, (©2000).

This product includes software developed by Object Oriented Concepts, Inc., Billerica, MA, USA (©2000).

This product includes free software developed by ImageMagick Studio LLC (©1999-2011).

This product includes software developed by Bob Withers.

This product includes software developed by Jean-Loup Gaily and Mark Adler.

This product includes software developed by Markus FXJ Oberhumer.

This product includes software developed by Guillaume Fihon.

This product includes QPDF software, developed by Jay Berkenbilt, copyright ©2005-2010, and distributed under version 2 of the OSI Artistic License (http://www.opensource.org/licenses/artistic-license-2.0.php).

This product includes JZlib software, Copyright © 2000-2011 ymnk, JCraft,Inc. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- The names of the authors may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED "AS IS" AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL JCRAFT, INC. OR ANY CONTRIBUTORS TO THIS SOFTWARE BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This product includes Apache Lucene software, distributed by the Apache Software Foundation under the Apache License, version 2.0.

This product includes Apache MINA software, distributed by the Apache Software Foundation under the Apache License, version 2.0.

This product includes OData4J software, distributed under the Apache License version 2.0.

This product includes software developed by the Visigoth Software Society (http://www.visigoths.org/).

This product includes software developed by Jeremy Ashkenas and DocumentCloud, and distributed under the MIT license. Copyright © 2010-2013 Jeremy Ashkenas, DocumentCloud.

This product includes software developed by Addy Osmani, and distributed under the MIT license. Copyright © 2012 Addy Osmani.

This product includes software developed by Charles Davison, and distributed under the MIT license. Copyright © 2013 Charles Davison.

This product includes software developed by The Dojo Foundation, and distributed under the MIT license. Copyright © 2010-2011, The Dojo Foundation.

This product includes gson software, distributed under the Apache License version 2.0. Copyright © 2008-2011 Google Inc.

This product includes software developed by Douglas Crockford, douglas@crockford.com.

This product includes ec2-tools software, copyright © 2008, Amazon Web Services, and licensed under the Amazon Software License. A copy of the License is located at http://aws.amazon.com/asl/ .

This product includes the ixgbevf Intel Gigabit Linux driver, Copyright © 1999 - 2012 Intel Corporation, and distributed under the GPLv2 license, as published by the Free Software Foundation.

This product includes Apache Ant software, distributed by the Apache Software Foundation under the Apache License, version 2.0.

This product includes libwebp software. Copyright © 2010, Google Inc. All rights reserved.

This product includes isc-dhcp software. Copyright © 2004-2013 by Internet Systems Consortium, Inc. ("ISC"); Copyright © 1995-2003 by Internet Software Consortium.

Permission to use, copy, modify, and/or distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies.

THE SOFTWARE IS PROVIDED "AS IS" AND ISC DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL ISC BE LIABLE FOR ANY SPECIAL, DIRECT, INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR

OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

This product includes jQuery Sparklines software, developed by Gareth Watts, and distributed under the new BSD license.

This product includes jsdifflib software, developed by Chas Emerick, and distributed under the BSD license.

This product includes winston software, copyright © 2010, by Charlie Robbins.

This product includes Q software developed by Kristopher Michael Kowal, and distributed under the MIT license. Copyright © 2009-2013 Kristopher Michael Kowal.

This product includes SlickGrid software developed by Michael Liebman, and distributed under the MIT license.

This product includes JCraft Jsch software developed by Atsuhiko Yamanaka, copyright © 2002-2012 Atsuhiko Yamanaka, JCraft, Inc. All rights reserved.

This product includes DP_DateExtensions software developed by Jim Davis, Copyright © 1996-2004, The Depressed Press of Boston (depressedpres.com). All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of the DEPRESSED PRESS OF BOSTON (DEPRESSEDPRESS.COM) nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

All code not authored by the Depressed Press is attributed (where possible) to its rightful owners/authors, used with permission and should be assumed to be under copyright restrictions as well.

This product includes Boost libraries, which are distributed under the Boost license (http://www.boost.org/LICENSE_1_0.txt).

This product includes Angular software developed by Google, Inc., http://angulargs.org, copyright © 2010-2012 Google, Inc., and distributed under the MIT license.

This product includes node.js software, copyright © Joyent, Inc. and other Node contributors. All rights reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

- The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

Legal Notices and Acknowledgments

This product includes node-uuid software, copyright © 2010-2012, Robert Kieffer, and distributed under the MIT license.

This product includes opencsv software, which is distributed under the Apache 2.0 license.

This product includes owasp-jave-encoder software, copyright © 2014, Jeff Ichnowski, and distributed under the New BSD license.

This product includes cookies software, copyright © 2014, Jed Schmidt, http://jed.is/, and distributed under the MIT license.

This product includes node-fastcgi software, copyright © 2013, Fabio Massaioli, and distributed under the MIT license.

This product includes socket.io software, copyright © 2013, Guillermo Rauch, and distributed under the MIT license.

This product includes node-querystring software, copyright © 2012. Irakli Gozalishvili. All rights reserved.

This product includes TinyRadius software, copyright © 1991, 1999 Free Software Foundation, Inc., and distributed under the GNU Lesser GPL version 2.1 license.

This product may include Intel SDD software subject to the following license; check your hardware specification for details.

1. LICENSE. This Software is licensed for use only in conjunction with Intel solid state drive (SSD) products. Use of the Software in conjunction with non-Intel SSD products is not licensed hereunder. Subject to the terms of this Agreement, Intel grants to You a nonexclusive, nontransferable, worldwide, fully paid-up license under Intel's copyrights to:

   • copy the Software onto a single computer or multiple computers for Your personal, noncommercial use; and
   • make appropriate back-up copies of the Software, for use in accordance with Section 1a) above.

   The Software may contain the software or other property of third party suppliers, some of which may be identified in, and licensed in accordance with, any enclosed "license.txt" file or other text or file.

   Except as expressly stated in this Agreement, no license or right is granted to You directly or by implication, inducement, estoppel or otherwise. Intel will have the right to inspect or have an independent auditor inspect Your relevant records to verify Your compliance with the terms and conditions of this Agreement.

2. RESTRICTIONS. You will not:

   a. copy, modify, rent, sell, distribute or transfer any part of the Software, and You agree to prevent unauthorized copying of the Software; and,
   b. reverse engineer, decompile, or disassemble the Software; and,
   c. sublicense or permit simultaneous use of the Software by more than one user; and,
   d. otherwise assign, sublicense, lease, or in any other way transfer or disclose Software to any third party, except as set forth herein; and,
   e. subject the Software, in whole or in part, to any license obligations of Open Source Software including without limitation combining or distributing the Software with Open Source Software in a manner that subjects the Software or any portion of the Software provided by Intel hereunder to any license obligations of such Open Source Software. "Open Source Software" means any software that requires as a condition of use, modification and/or distribution of such software that such software or other software incorporated into, derived from or distributed with such software:

      a. be disclosed or distributed in source code form; or
      b. be licensed by the user to third parties for the purpose of making and/or distributing derivative works; or
      c. be redistributable at no charge.

Open Source Software includes, without limitation, software licensed or distributed under any of the following licenses or distribution models, or licenses or distribution models substantially similar to any of the following:

a. GNU's General Public License (GPL) or Lesser/Library GPL (LGPL),

b. the Artistic License (e.g., PERL),

c. the Mozilla Public License,

d. the Netscape Public License,

e. the Sun Community Source License (SCSL),

f. vi) the Sun Industry Source License (SISL),

g. vii) the Apache Software license, and

h. viii) the Common Public License (CPL).

3. OWNERSHIP OF SOFTWARE AND COPYRIGHTS. Title to all copies of the Software remains with Intel or its suppliers. The Software is copyrighted and protected by the laws of the United States and other countries, and international treaty provisions. You may not remove any copyright notices from the Software. Intel may make changes to the Software, or to materials referenced therein, at any time and without notice, but is not obligated to support or update the Software. Except as otherwise expressly provided, Intel grants no express or implied right or license under Intel patents, copyrights, trademarks, or other intellectual property rights.

4. Entire Agreement. This Agreement contains the complete and exclusive statement of the agreement between You and Intel and supersedes all proposals, oral or written, and all other communications relating to the subject matter of this Agreement. Only a written instrument duly executed by authorized representatives of Intel and You may modify this Agreement.

5. LIMITED MEDIA WARRANTY. If the Software has been delivered by Intel on physical media, Intel warrants the media to be free from material physical defects for a period of ninety (90) days after delivery by Intel. If such a defect is found, return the media to Intel for replacement or alternate delivery of the Software as Intel may select.

6. EXCLUSION OF OTHER WARRANTIES. EXCEPT AS PROVIDED ABOVE, THE SOFTWARE IS PROVIDED "AS IS" WITHOUT ANY EXPRESS OR IMPLIED WARRANTY OF ANY KIND, INCLUDING WARRANTIES OF MERCHANTABILITY, NONINFRINGEMENT, OR FITNESS FOR A PARTICULAR PURPOSE. Intel does not warrant or assume responsibility for any errors, the accuracy or completeness of any information, text, graphics, links or other materials contained within the Software.

7. LIMITATION OF LIABILITY. IN NO EVENT WILL INTEL OR ITS SUPPLIERS BE LIABLE FOR ANY DAMAGES WHATSOEVER (INCLUDING, WITHOUT LIMITATION, LOST PROFITS, BUSINESS INTERRUPTION OR LOST INFORMATION) ARISING OUT OF THE USE OF OR INABILITY TO USE THE SOFTWARE, EVEN IF INTEL HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. SOME JURISDICTIONS PROHIBIT EXCLUSION OR LIMITATION OF LIABILITY FOR IMPLIED WARRANTIES OR CONSEQUENTIAL OR INCIDENTAL DAMAGES, SO THE ABOVE LIMITATION MAY NOT APPLY TO YOU. YOU MAY ALSO HAVE OTHER LEGAL RIGHTS THAT VARY FROM JURISDICTION TO JURISDICTION.

8. TERMINATION OF THIS AGREEMENT. Intel may terminate this Agreement at any time if You violate its terms. Upon termination, You will immediately destroy the Software or return all copies of the Software to Intel.

9. APPLICABLE LAWS. Claims arising under this Agreement will be governed by the laws of Delaware, excluding its principles of conflict of laws and the United Nations Convention on Contracts for the Sale of Goods. You may not export the Software in violation of applicable export laws and regulations. Intel is not obligated under any other agreements unless they are in writing and signed by an authorized representative of Intel.

10. GOVERNMENT RESTRICTED RIGHTS. The Software is provided with "RESTRICTED RIGHTS." Use, duplication, or disclosure by the Government is subject to restrictions as set forth in FAR52.227-14 and DFAR252.227-7013 et seq. or their successors. Use of the Software by the Government constitutes

Index

Index

general properties
host IP address *33*
host name *33*
management port and TMM *33*
management port configuration *33*
time zone *33*

H

heartbeat failure *37*
hotfix
importing *24*
installing *24*
HTTP compression strategies
about *41*

I

IP geolocation data *7*
ISO
importing *23*
installing *23*

L

last hop
automatically mapping *9*
Layer 2 forwarding table
record cache time setting *9*
licenses
activating automatically *27*
activating manually *27*
reactivating *28*
licensing
about *27*
local traffic
properties *9*

M

MAC address
sharing among VLANs *9*
maintenance mode
connections with *9*
MTU discovery
effect on TCP packet fragmentation *9*

N

Network Time Protocol (NTP)
8
configuring the time server list *8*
NTP, See Network Time Protocol (NTP).

P

password
administrative account *34*
platform properties
about *33*
configuring *34*

R

reaper high-water mark
described *9*
reaper low-water mark
described *9*
redundant device properties
about *34*
resource provisioning
about *29*
running configuration
about *11*

S

SCF
See also single configuration file (SCF).
propagating system data *15*
See also single configuration file (SCF).
single configuration file (SCF)
about *11*
copying *12*
creating *12–13*
loading *12*
managing using tmsh *11*
saving *12–13*
tmsh commands *13*
SNATs
forwarding packets from *9*
software image
importing *23*
installing *23*
software management
about *21*, *23*
SSH access configuration *34*
stored configuration
about *11*
SYN cookies
activation for TCP connections *9*
system configuration data
about *11*
managing *11*
running *11*
stored *11*
system fail-safe
about *37*
configuring *37*
system provisioning
using the Configuration utility *29*
using the tmsh *29*

T

tmsh
creating an archive *19*
deleting an archive *20*
loading archive data *20*
managing archives *19*
restoring archive data *20*
restoring archives *20*
saving an archive *19*
viewing a list of archives *19*

Index