CA AppLogic®

Release Notes

3.5

19th Edition: Hotfix hf9042 (July 1, 2013)

Contact CA Technologies

Contact CA Support

For your convenience, CA Technologies provides one site where you can access the information that you need for your Home Office, Small Business, and Enterprise CA Technologies products. At http://ca.com/support, you can access the following resources:

- Online and telephone contact information for technical assistance and customer services
- Information about user communities and forums
- Product and documentation downloads
- CA Support policies and guidelines
- Other helpful resources appropriate for your product

Providing Feedback About Product Documentation

If you have comments or questions about CA Technologies product documentation, you can send a message to techpubs@ca.com.

To provide feedback about CA Technologies product documentation, complete our short customer survey which is available on the CA Support website at http://ca.com/docs.

Documentation Changes

The following documentation updates have been made since the last release of this documentation:

19th Edition: Hotfix hf9042 (July 1, 2013)

- Hotfixes for CA AppLogic® > hf9042 (see page 99): Resolves an issue in CA AppLogic® 3.5.19 and 3.7.14 where applications containing a swap volume cannot be migrated (SCR 9042).

18th Edition: Hotfix df8975 (June 14, 2013)

- Hotfixes for CA AppLogic® > df8975 (see page 100): Resolves several device support, security and stability issues (for new grids).

17th Edition: Hotfix hf7596 (June 11, 2013)

- Hotfixes for CA AppLogic® > hf7596 (see page 105): Resolves an issue in CA AppLogic® 3.5.19 where servers with very large data volumes can fail on reboot or upgrade due to a filesystem check running the server out of memory (SCR 7596).

16th Edition: Hotfix hf8096 (June 7, 2013)

- Hotfixes for CA AppLogic® > hf8096 (see page 105): Resolves the following: SCR 8096: Product: Xen: HVM guest freezes during heavy I/O. Also resolves Xen Security Advisories.

15th Edition: Hotfix hf8315 (May 23, 2013)

- Hotfixes for CA AppLogic® > hf8315 (see page 108): Resolves SCR 8315 - Product: SAN ESX grid upgrade failure to CA AppLogic® 3.7.x

14th Edition: Hotfix hf8114 (February 14th, 2013)

- Hotfixes for CA AppLogic® > hf8114 (see page 108): Resolves the following: SCR 8114 - Product: lost access to the grid controller while managing a windows volume and copying a file to it from the impex volume. SCR 7673 - Xen netback driver crashes when receiving large packets.

13th Edition: Hotfix hf7781 (February 9th, 2013)

- Hotfixes for CA AppLogic® > hf7781 (see page 108): Resolves SCR 7781 - heartbeat logs fill filesystem on dom0 (which may cause a server to fail to boot).

12th Edition: Hotfix e7785 (February 5th, 2013)

- Hotfixes for CA AppLogic® > e7785: Provides an enhancement to CA AppLogic® that upgrades the Web Services API application template to version 1.0.20 (r20).

11th Edition: hf8017 (January 17, 2013)

- Hotfixes for CA AppLogic® > hf8017 (see page 109): Resolves Xen Security Advisories (XSAs) 20 - 31.

10th Edition: BFC 3.5.2 and hf8003 (January 11, 2013)

- Backbone Fabric Controller (BFC) New Features > VLAN 0 Support (BFC 3.5.2) (see page 20): How to add VLAN 0 to a network and select it on a grid.
- Backbone Fabric Controller (BFC) New Features > Performance Enhancements in the BFC 3.5.2 (see page 23): Describes performance optimizations with this release of the BFC.
- Key Bug Fixes (see page 78): Updated for BFC defects resolved with BFC 3.5.2.
- Hotfixes for CA AppLogic® > hf8003 (see page 119): Resolves SCR 7673 - Xen netback driver crashes when receiving large packets.

9th Edition: Hotfix hf7906 (December 13, 2012)

- Hotfixes for CA AppLogic® > hf7906 (see page 119): Resolves an issue where applications open slowly (~2 minutes) in the infrastructure editor when opened by a user who is a member of a large number of groups (SCR 7906).

8th Edition: Hotfix hf7791 (November 21, 2012)

- Hotfixes for CA AppLogic® > hf7791 (see page 111): Upgrades the HP Smart Array Controller driver hpsa to the latest version 3.1.0-7.

7th Edition: Hotfix hf7759 (November 9, 2012)

- Hotfixes for CA AppLogic® > hf7759 (see page 112): Resolves Xen Security Advisory XSA-25 http://seclists.org/oss-sec/2012/q4/141, where an Out-of-memory error can occur due to a malicious kernel/ramdisk.

6th Edition: Hotfix e7748 (November 2, 2012)

- Hotfixes for CA AppLogic® > e7748 (see page 120): Provides an enhancement to CA AppLogic® where it automatically restarts an appliance when CA AppLogic® loses connection to the VM agent running within the appliance.

5th Edition: Hotfixes hf7720 and hf7330 (October 25, 2012)

- Known Issues > Hotfixes for CA AppLogic® > hf7720 (see page 120): Resolves an issue where importing a class/catalog from a URL does not work (SCR 7720).
- Known Issues > Hotfixes for CA AppLogic® > hf7330 (see page 120): Resolves an issue where migrating applications is slow between grids that are geographically separated (SCR 7330).

4th Edition: BFC 3.5.1 and Hotfixes hf7694 and df7588 (October 12, 2012)

- Backbone Fabric Controller (BFC) New Features > Support for CentOS 5.8 (3.5.1): Improvements in packaging and distribution of the Bare Metal Installer using CentOS 5.8, and corrections to the core 3.5 documentation noting the required CentOS version.
- Backbone Fabric Controller (BFC) New Features > Bare Metal Installer (3.5.1 update) (see page 21): The 3.5.1 release of the Bare Metal Installer includes CentOS 5.8.
- Installation Considerations > Installation, Upgrade, and Migration > CentOS 5.8 and BFC 3.5.1 Upgrade: How to upgrade CentOS before installing BFC 3.5.1.1969
- Known Issues > Hotfixes for CA AppLogic® > df7588 (see page 115): Resolves Xen device support issues.
- Known Issues > Hotfixes for CA AppLogic® > hf7694 (see page 112): Resolves Xen device support issues.
- Key Bug Fixes: Updated for BFC defects resolved with 3.5.1.

3rd Edition: Hotfix hf7611 (September 18, 2012)

- Known Issues > Hotfixes for CA AppLogic® > hf7611 (see page 121): Resolves Xen Security Advisories.

2nd Edition: Hotfix hf6169 and hf7516 (August 24, 2012)

- Known Issues > Hotfixes for CA AppLogic® > hf6169: Fixes an issue where Windows appliances fail to start after their hostname has been changed.
- Known Issues > Hotfixes for CA AppLogic® > hf7516 (see page 122): Fixes an issue where the link state for an Emulex 10G NIC is wrongly reported as unknown or down due to the bug in the Emulex NIC driver be2net.

Contents

Chapter 1: New in This Release

Version 3.5.19 - August 13, 2012

Note: This release is the official CA AppLogic® General Availability release and is available immediately to partners and customers who participate in the General Availability program. The current production release of CA AppLogic® is 3.5.19.

This section describes new features and enhancements included in CA AppLogic® 3.5.

This section contains the following topics:

Storage Area Network (SAN) Support Over NFS

This release lets you connect your existing SAN to your CA AppLogic® grid using the IP-based Network File System (NFS) protocol. Multiple grids running within the same backbone can share a SAN. For each grid, CA AppLogic® uses a unique directory on the SAN where all of the volumes for the grid are stored.

- This new functionality maintains the current interfaces, including the GUI, CLI and API, for activities such as volume creation, removal, and resizing.
- When you create a grid, you enter the IP address and the name of the NFS share that the SAN exports.
- When you configure a CA AppLogic® grid to use an external SAN, CA AppLogic® stores all of the application and appliance volumes on the SAN by default, including the grid controller and system volumes. You can use the *vol migrate* command to migrate volumes between the internal SAN and the external SAN.
- Many of the dialogs and windows in CA AppLogic®, including the Create Volume dialog, have a new drop-down to select local or SAN storage.
  - If SAN is selected, the volume is created and stored on the SAN.
  - If local is selected, the volume is created within the local CA AppLogic® storage. This behavior matches the way CA AppLogic® behaves in all previous releases.

- CA AppLogic® does not explicitly mirror volumes that are created on the SAN. CA AppLogic® relies on the mirroring and resiliency of the SAN for storage integrity.

Note: For more information about this feature, see the *User Interface Reference Guide*.

Dynamic Canvas: Status and Control

This feature lets you monitor component state changes within an application, and start and stop these components, all from the infrastructure editor.

- You can see the current state of components within an application, including the starting, stopping, running, and error states of each component.
- Each component icon contains a new control mechanism that lets you start and stop the component within an application.
- A new drop-down dialog displays progress and status during application and component start/stop operations.

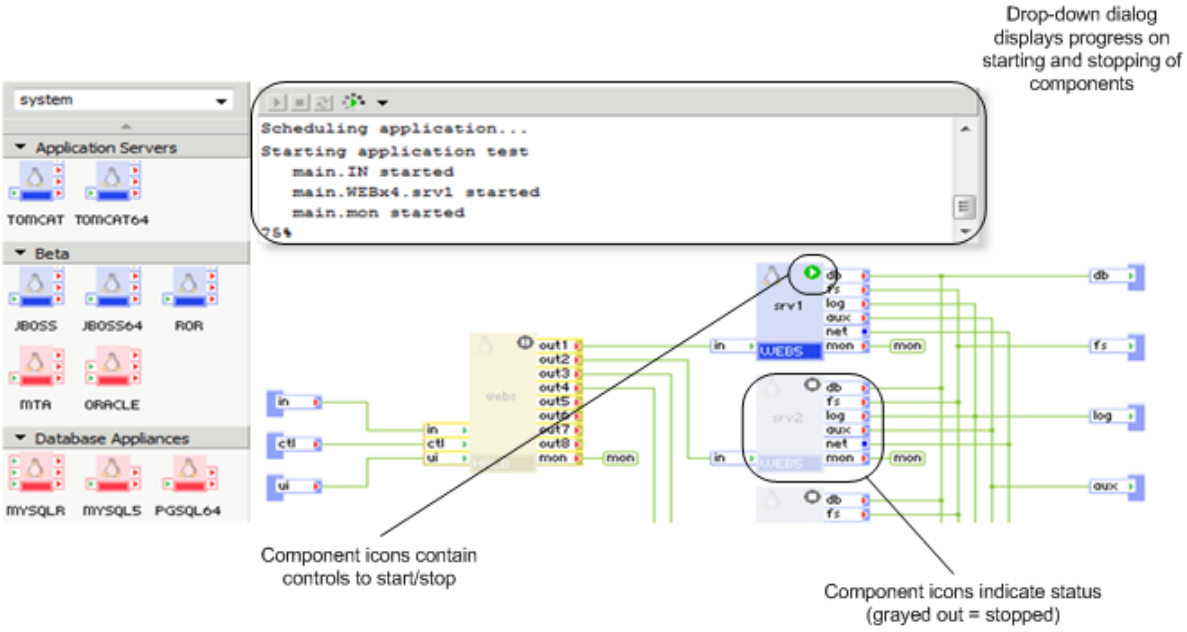

Note: For more information about this feature, see Common Operations in the *Grid User Guide*, and Editing Applications in the *User Interface Reference Guide*.

Quotas

This release introduces a new method to control the consumption of resources using predefined quotas.

- This new functionality is available in both the CLI and Web API.
- You can assign quotas to users or groups for the following resources: CPU, Memory, Disk, and Bandwidth on the internal network.

  Note: Quotas are disabled by default and do not constrain resources for users.
- After you configure the quotas, CA AppLogic® monitors the quotas at the time of resource allocation for applications, catalog, and volumes.

Note: For more information about this feature, see Quotas in the *Grid Administration Guide*.

Multiple External Interfaces

In previous CA AppLogic® releases, appliances supported no more than one external (raw) interface. As a result, they were also limited to using a single VLAN. The new Multiple External Interface (MEI) feature enables appliances to use multiple external interfaces and consequently multiple VLANs.

External interfaces are now visible on the appliance boundary, just like with regular input and output terminals. In previous releases external interfaces had no visual representation in the infrastructure editor. External interfaces are now connected to raw interfaces that are defined on the application boundary. These raw interfaces represent an application's connections to the outside world and are now explicitly shown in the diagram, just like regular terminal connections between appliances. In addition, the new APK that is shipped with the release now automatically configures the external interfaces of an appliance.

There is a new **Interfaces** tab in the Application Configuration editor that shows the application's raw interfaces and their settings. Network (VLAN) and IP addresses are the only configurable settings selected by the user for each raw interface. Netmask and gateway settings are automatically derived from the grid's configuration, selecting settings that correspond to the chosen network and IP address. The CLI and Web services API have also been extended to allow the configuration of an application's IP settings, such as app config and app provision.

The BFC is used to define the VLANs and IP ranges within each VLAN. These settings are propagated to the CA AppLogic® grid and are displayed on the grid's dashboard. These are the settings that may be configured on the raw (external) interfaces exposed by an application.

With the addition of the multiple external interface feature, backwards compatibility is fully maintained for existing appliances that use the old style external interfaces. Older appliances that use a single external interface receive their network/IP configuration through properties. However, appliances that use the new style external interfaces are not compatible with older versions of CA AppLogic® that do not support the multiple external interface feature (releases prior to CA AppLogic® 3.5).

Important! The default catalog and sample applications have not been updated to support multiple external interfaces. If you want to update one of these appliances to use multiple external interfaces, manually change the appliance script and install the new APK (see page 15) from the CA AppLogic® 3.5 General Availability release. The catalog and sample applications will be updated to support multiple external interfaces in a future release.

Note: For more information, see Define Multiple Interfaces in the *Appliance Developer Guide*.

Advanced Application Start/Restart Options in the User Interface

The CA AppLogic® GUI now lets you select various advanced start options when starting and restarting applications. In previous releases, there were no start/restart options available to select from the GUI, and the only way to specify options for starting applications was through the CLI or web API.

The following options are now available in the GUI:

- Scheduling mode: Pack Servers (default), Round robin and Round robin with random start.
- Debug Mode
- CPU Capping

Note: For more information, see the *Command Line Shell Reference Guide* and the *Grid User Guide*.

Volume Repair Speed Control

You can now control the volume repair speed used by the grid using the grid set command (grid set vol_repair_speed). This enables the degraded volumes within the grid to be repaired faster, which speeds up grid recovery. By default, the volume repair speed is only 10MB/sec, which can be very slow depending upon volume size.

Note: Exercise caution when changing the volume repair speed because this can impact the performance of running applications. The volume repair speed can only be changed by grid maintainers.

Other CA AppLogic® Enhancements

This section contains the following topics:

New Windows APK

The new Windows Appliance Kit (APK) provides an easier installation and user-friendly interface for configuring Windows-based appliances. Support for the localized versions of the Windows OS is planned for a future release.

Note: The *Appliance Developer Guide* reflects the previous Windows APK. However, you can access an updated version of the affected topics from this guide for the new Windows APK (APP_3_5_readme_ENU.html).

If you participated in the beta program, review the upgrade considerations (see page 33).

Support Utility (3tcollect)

This release includes the 3tcollect utility. The *3tcollect utility* is a command-line tool that you can execute on the Backbone Fabric Controller (BFC) or the grid controller to collect log and configuration data.

Note: For more information about this feature, see Support Utility in the *Grid Administration Guide*.

Hidden Property Values

This release includes a hidden attribute for parameter fields to secure property values, such as passwords. Using the hidden attribute for a property hides the property value so that regular users cannot see the value.

Note: For more information about this feature, see the *User Interface Reference Guide*.

Enhanced Hardware Compatibility

This release supports enhanced hardware compatibility by including the latest stable version of the Linux kernel 3.2.2, Xen 4.1.2, and also supporting ESX 4.1 update 2.

Cloud Commons Marketplace Integration

This release includes a Cloud Commons integration between the Marketplace servers and the customer grids. This integration facilitates the delivery of purchased items from the Marketplace.

- After you purchase a product, this integration transmits the product to your grid.
- The CA AppLogic® grid adds the product to private catalogs.

Documentation Enhancements

This section describes enhancements to the CA AppLogic® documentation.

Note: This topic only describes general enhancements to the documentation. For details about the technical changes to the documentation, see the Documentation Changes topics at the beginning of each guide.

All-in-One HTML

The previous bookshelf contained links that opened individual guides in an HTML window that only had access to that guide. If you wanted to access another guide in the bookshelf, you had to click the Back to Bookshelf link at the top of the TOC pane.

For this release, the links open an HTML window that contains all of the guides, greatly improving navigation between guides.

End-to-End Bookshelf

The previous bookshelf primarily contained links to the guides. For this release, we are introducing an End-to-End bookshelf that also contains links to articles, social forums, and videos, among other things.

Updated Bookshelf Categories

We have changed the names of the bookshelf categories to reflect their content better, and we have also moved some of the guides into different categories. These categories are also included in the all-in-one HTML window.

New Guides

We have created the following new guides:

Grid Administrator Guide

This guide is comprised of the following content:

- Content previously in the *Role-Based Access Control (RBAC) User Guide*
- New content to describe the Quota functionality added in this release.

Advanced Maintenance Guide

This guide contains content previously in the *Backbone Fabric Controller (BFC) User Guide*. We moved this content because it is not related to BFC.

Backbone Fabric Controller (BFC) API Guide

This guide describes how to use the new BFC web API.

Documentation Changes Topics

We now include a Documentation Changes topic at the beginning of each guide that lists changes related to the marquee features described in this Release Notes, and other important changes.

Glossary in *Overview Guide*

We now include a glossary in the *Overview Guide* to define some of the CA AppLogic® terms that are used throughout the documentation.

Backbone Fabric Controller (BFC) New Features

This section contains the following topics:

Support for CentOS 5.8

The 3.5.2 version of the Backbone Fabric Controller supports and requires CentOS 5.8.

- The CentOS 5.8 image included in the Bare Metal Installer has been hardened to reduce the number of delivered packages.
- The number of packages that are delivered with the Bare Metal image have been reduced to the minimum number of packages the BFC requires for its operations. This results in a smaller image, and significantly reduces the time that is required to download the file. The smaller number of packages also results in faster installation speed overall.

Several topics in the 3.5 core documentation refer to the previous requirement for CentOS 5.5. Note the following documentation corrections:

Correction to the BFC control node requirements:

For BFC 3.5.2, unless you plan to use the Bare Metal Installer feature, install CentOS 5.8 in advance.

Correction to procedure for installing the BFC on an existing CentOS system:

Begin by installing a standard package configuration of CentOS 5.8 to the server where you plan to install or upgrade the Backbone Fabric Controller.

Correction to system requirements for the Bare Metal ISO tool:

Before you install the Bare Metal ISO tool, make sure that your operating system is CentOS 5.8.

BFC Application Programming Interface (API)

The BFC Application Programming Interface (API) enables you to work with Create, Remove, Update, and Delete functionality for grids using a Web Services-based API. The API lets external systems initiate functions within BFC using a programmatic interface.

The interface is based on the web protocol and provides a means to send and receive results to the system using the HTTPS protocol. Results of executed commands are returned in a standard JSON format.

Note: For more information, see the Backbone Fabric Controller (BFC) API Reference Guide.

Grid Node Selection Using Tags

This release provides an enhanced mechanism for the selection of server resources during the grid creation process.

- You work with new dialogs and management interfaces to tag servers in the available server pool using user-defined criteria.
- During the Grid Create workflow, you select one or more tags to determine the pool of server resources for building a new grid. In addition, you can select the individual servers to create a grid.

You can manage servers using tags in a new Tag tab on the BFC Administration page. Grid and Server pages are enhanced to let you assign tags and display their usage.

Configure VLAN Ranges

New features have been added to support configuring VLAN ranges for use by your applications in the grid:

- You can use options in the Administration tab -> Networks -> External Networks page to add and configure VLANs for a network.
- A new Network Resources page is displayed in the grid create wizard where you can specify IP ranges to allocate for the grid.
- You can also edit VLAN information for a grid in the grid properties page, Network tab.

VLAN 0 Support (BFC 3.5.2)

The way you can work with tagged and untagged networks in the same grid has changed with the BFC 3.5.2. In the BFC 3.5.2, there is a provision to associate VLAN ID 0 with a grid. This is useful in cases where you want your grid controller to be on a public network (no VLAN tagging) while the appliances on the grid are still assigned to their respective VLANs. For such a case, you can associate VLAN 0 with the controller and associate different VLAN(s) to the application IPs.

In the BFC 3.5.0 and 3.5.1, VLAN 0 was not a valid VLAN ID. In the BFC 3.5.2 and later, you can add VLAN 0 to a network in the Administration page, Networks tab, VLANs option. You then can then select it when you create a grid, or update grid properties. There are some restrictions for selecting VLAN 0 on a grid, and some considerations for upgrade described below.

Consider the following information before you add VLAN 0:

- You can add VLAN 0 from the Grid Create dialog or for an existing grid on the Grid Properties, Networks page. Note that certain constraints apply when working with VLAN 0.

  The BFC treats VLAN 0 as a tagged VLAN. If the grid already has tagged VLANs, you can add VLAN 0 to the grid. If the grid already has untagged VLANs, you cannot add VLAN 0.

- When you add VLAN 0, a warning message displays on the following dialogs:
  - Administration tab, Networks tab, VLANs, Add VLAN
  - Grid Create/Properties dialog

  Warning: Using IP addresses on VLAN 0 can leave the system vulnerable to a VLAN hopping attack as well as duplication of addresses.

  Note: If you click yes, you add the range and the message dialog closes. If you click no, the message dialog closes and the Add a VLAN dialog remains. To avoid oversubscription of addresses, ensure that no devices external to the BFC are using the address range that you indicate for VLAN 0.

- If the BFC detects addresses that use VLAN 0 during an upgrade, a warning message displays:

  This BFC is using IP addresses on VLAN 0. Using IP addresses on VLAN 0 can leave the system vulnerable to a VLAN hopping attack as well as duplication of addresses.

  Note: If grids have VLANs configured in the BFC but not actually in use on the grid, you can remove those unused VLANs to reduce the time required to upgrade the BFC. You can select to cancel the upgrade and remove the unused VLANs, or you can proceed with the upgrade.

Bare Metal Installer

This release provides a repackaging of the BFC into a full distribution that includes the operating system (CentOS 5.8). This repackaging improves installation efficiency and accuracy.

This new product format provides a means to install both the operating system and the product from the *bare metal* (from the minimum BFC system requirements). The only requirement is a system that meets the minimum system requirements for the BFC.

The following ISO images are distributed:

- Operating system
- BFC installer

In addition, the Bare Metal ISO tool is provided to allow you to create a single ISO image from the two distributed ISO images. You can modify a sample template configuration file to customize the parameters for use during both attended and unattended installations, on local or remote hardware. You can also use the single ISO install image for quick BFC restore, and also to create variations of the installation (for example, to specify different IP addresses).

Notes:

- For an overview of new features, see the topic Bare Metal Installer.
- For more information on working with the Bare Metal ISO tool, see the topic Create an Install Image Using the Bare Metal ISO Tool.

Network Configuration Diagnostics

The Hardware tab of the Server Properties Page in the BFC has been enhanced to display more information about NICs found during server discovery by the BFC.

In previous releases, this page only displayed information about the backbone and external networks. Now it includes detailed information for all discovered servers, including:

- MAC address
- Network type
- IP Address
- Interface type
- NIC speed
- Switch ID
- Switch type

- Switch protocol
- STP port

Reviewing this diagnostic data helps you analyze potential server issues, for example:

- When a grid fails due to a network problem.
- When a system is quarantined due to network problems detected during discovery.

Additional information helps you diagnose configuration issues, for example, to check if two NICs use the same switch ID. The information also helps you easily identify your backbone and external networks.

You can customize the display of this information using a drop-down menu that appears when you click to the right of any column name.

Notes:

- For more information, see Review Network Configuration Data in the *Backbone Fabric Controller (BFC) User Guide*.
- Due to the additional checks done as part of this new functionality, some non-compliant configurations previously allowed are now disallowed. Although these configurations were allowed in the past, they did not function reliably.

  The following configurations are now disallowed:

  - A server with the STP switch protocol detected on the backbone NIC
  - A server with the STP switch protocol detected on the external NIC, unless the port in question is configured as “portfast” or “edge port”
  - A server with a backbone NIC with a speed of less than 1Gbps

Force Manual Power

A new option has been added to the BFC to enable you to address power controller issues in a server by forcing it to be manually powered. You can undo the action to restore the server status to its original state (usually IPMI power-controlled).

You can access this feature on the Server Overview page when you select Force Manual Power or Undo Force Manual Power from the Server Actions drop-down list.

Note: For more information, see Force Manual Power in the *Backbone Fabric Controller (BFC) User Guide*.

Performance Enhancements in the BFC 3.5.2

Optimizations introduced with the BFC 3.5.2 have resulted in the following scalability and performance improvements:

- Creation time for VLANs significantly improved. The test environment included creating up to 2000 VLANs.
- Time taken to update VLANS to the grid has improved.
- The display of VLAN lists in the user interface has been optimized. VLAN lists display in the Grid Properties page, and in the VLANs dialog you open from the Administration page, Networks tab.
- With 2000 VLANS in all grids, BFC service start timing has improved significantly.
- Deletion time for VLANs has significantly improved.
- Optimizations include resolving outstanding bugs. See the list of key bug fixes (see page 78) for more information.

Chapter 2: Installation Considerations

Before you begin your installation, check the CA AppLogic® home page on the CA Support Online website for the latest news.

You can either install CA AppLogic® 3.5 as a new installation of CA AppLogic®, or upgrade from Version 3.0, or Release 3.1.

This section contains the following topics:

Installation, Upgrade, and Migration

This section describes how to install, upgrade, and migrate CA AppLogic® 3.5.

This section contains the following topics:

CentOS 5.8 and BFC 3.5.2 Upgrade

If you want to do a fresh install of the BFC 3.5.2, you must confirm that CentOS 5.8 is installed on your system. (**Note:** If you plan to use the BFC Bare Metal Installation feature, the CentOS 5.8 installation is done automatically as part of the Bare Metal Installation process).

If you want to upgrade the BFC from 3.5.0 (or earlier) to 3.5.2, you must upgrade to CentOS 5.8. Use the following process to upgrade from CentOS 5.5 to CentOS 5.8.

Follow these steps to upgrade CentOS with a locally configured yum repository:

1. Download the CentOS 5.8 images from www.centos.org. The download should contain two DVD iso files (DVD1 and DVD2). Copy the downloaded files to the BFC machine.
2. Shut down the BFC service (service bfc stop)
3. On the BFC machine where you want to upgrade CentOS from 5.5 to 5.8, run the following command to verify that CentOS 5.5 is installed and ready for upgrade:

   rpm -import /mnt/CentOS/5.5/RPM-GPG-KEY-CentOS-5

   Note: This step assumes your current CentOS 5.5 yum configuration points to /mnt/CentOS/5.5/ as the yum repository.
4. On the BFC machine, mount the 2 DVD iso files separately on two directories:

   # mount –r –o loop CentOS5.8-DVD1.iso <dir1>

   # mount –r –o loop CentOS5.8-DVD2.iso <dir2>
5. Create a directory /mnt/centos58/ and copy the contents of the two directories (dir1 and dir2) to /mnt/centos58/.
6. Update the "baseurl=" parameter of the /etc/yum.conf:

   baseurl=file:///mnt/centos58/
7. Run the following YUM command:

   yum update
8. Reboot the BFC node for the updated packages (e.g. the kernel package) to get persistent.
9. Run the following YUM command to list the CentOS 5.8 packages:

   yum list.

Follow these steps to upgrade CentOS with the online yum repository

1. Verify that /etc/yum.repos.d/ directory exists.
2. Check whether the BFC machine has internet access (use the *wget* command to get any file from internet).

3. Run *yum update* to update the current installed CentOS 5.5 packages to CentOS 5.8 version).
4. Reboot the BFC node for the updated packages (e.g. the kernel package) to get persistent.
5. Run *yum list* to list the CentOS 5.8 packages.

   Note: For further details refer to your CentOS manuals for yum configuration.

Upgrade From 3.0 or 3.1 - Check Version Download Location

The following note only applies to users upgrading from version 3.0 or 3.1 to 3.5.

Before you upgrade CA AppLogic® 3.0 or 3.1 to CA AppLogic® 3.5, confirm that your version download directory is not set to /opt/bfc. The uninstaller that is provided with version 3.0 and 3.1 removes all files from the BFC install directory. If the 3.5 upgrade does not succeed for some reason and you use the 3.0/3.1 fallback to recover, all files in opt/bfc are removed.

This is not an issue if you do not use opt/bfc as your version download directory, and is not an issue for new 3.5 installations. Follow these steps to check your version download location.

1. Start the BFC, open the Administration page, and click the Versions tab.
2. Check the CA AppLogic® Version Downloads Directory field. If the directory path includes /opt/bfc, change it to another location. The directory must have read/write permissions for the bfcadmin user.
3. Move the current version download files to the new directory location. For example, enter the following command on the BFC machine to move downloaded files from /opt/bfc/downloads to /var/applogic/downloads:

   ```
   mv /opt/bfc/downloads /var/applogic/downloads
   ```

BFC Replaces ALD

ALD is no longer used to install and upgrade grids. Taking the place of ALD is the Backbone Fabric Controller (BFC). BFC contains a simple-to-use web-based GUI application that you use to create and manage all of your CA AppLogic® grids within a single backbone. BFC automatically downloads the latest CA AppLogic® releases and hotfixes.

The BFC documentation includes information about how to download/install BFC and how to use it to manage your CA AppLogic® grids.

Upgrading Existing Grids

This section contains the following topics:

Upgrades from existing CA AppLogic® Version 3.0, 3.1, 3.5, and 3.7 grids to this release (3.5) are fully supported for Xen-based grids. For ESX-based grids, upgrade from CA AppLogic® 3.0, 3.1, 3.5, and 3.7 to this release is *not* supported.

Upgrades are not supported from any CA AppLogic® release prior to 3.0, such as 2.9 and 2.8. To migrate your older CA AppLogic® grid to the latest release, you must install CA AppLogic® 3.5 and migrate your custom applications and catalogs from your old grid to your new 3.5 grid.

Note: To upgrade Xen-based 3.x grids, refer to the BFC documentation.

CA AppLogic® 3.5 does support appliances and applications that were created with older CA AppLogic® releases. Depending upon the type of appliance and the kind of hypervisor that is needed, you may need to update the appliance before using it on a 3.5 grid.

Note: Solaris-based appliances are no longer supported in CA AppLogic® 3.5.

Determine Changes for Your Applications

Follow the guidelines below to determine what kind of modifications (if any) are needed to make your applications run on a 3.5 grid:

- If you move appliances from a 3.0/3.1/3.5 grid to a 3.5 grid (Xen or ESX-based), everything works as-is without any modifications.
- If moving appliances from a 2.9 grid (or prior version) to a 3.5 Xen-based grid, everything will work as-is w/o any modifications
- If you move Windows-based appliances from a 2.9 grid (or prior) to a 3.5 ESX-based grid, the descriptors need to be updated to the new 3.x format (see below), the latest APK (Appliance Kit) needs to be re-installed in the appliance and vmware-tools should be installed in the appliance too (vmware-tools can be downloaded from VMware's website)
- If you move Linux-based appliances from a 2.9 grid (or prior) to a 3.5 ESX-based grid, the descriptors need to be updated to the new 3.x format (see below) and the boot disk needs to be updated to support ESX (see below)
- If any appliance has more than 10 terminals and the intent is to run the appliance on an ESX-based grid, the descriptors need to be updated to use the new 3.x format (see below) and the latest APK (Appliance Kit) needs to be re-installed in the appliance

Update Appliance Descriptors

Follow these steps to update an appliance descriptor to use the new 3.x format.

Note: Perform this procedure if you have appliances from older CA AppLogic® releases that you want to run on the ESX hypervisor.

Follow these steps:

1. Import the appliance into the 3.x grid using the class import command (or import your custom catalogs or applications as needed). Import anything that contains appliances that you want to update to use the new 3.x ADL descriptor format.
2. If the appliance is not a singleton, create a new application, drag an instance of your appliance into the application and branch the appliance. If the appliance is a singleton, edit the application in the infrastructure editor.
3. Modify the appliance class and under the Advanced section of the General tab, select the appropriate virtualization modes based on the hypervisors are compatible with the appliance (more than one mode may be selected as long as the appliance supports all of the selected modes).
4. Select the Options button next to the virtualization modes and do the following:
   a. Select the VMware virtualization mode.
   b. In the options field, add a new setting named esx_os_name where the value is one of the following:
      - Microsoft Windows Server 2003, Datacenter Edition (64-bit): winNetDatacenter-64
      - Microsoft Windows Server 2003, Enterprise Edition (64-bit): winNetEnterprise-64
      - Microsoft Windows Server 2003, Standard Edition (64-bit): winNetStandard-64
      - Microsoft Windows Server 2003, Datacenter Edition (32-bit): winNetDatacenter
      - Microsoft Windows Server 2003, Enterprise Edition (32-bit): winNetEnterprise
      - Microsoft Windows Server 2003, Standard Edition (32-bit): winNetStandard
      - Microsoft Windows Server 2003, Web Edition: winNetWeb
      - Microsoft Windows Server 2008 R2 (64-bit): windows7srv-64
      - Microsoft Windows Server 2008 (32-bit): longhorn
      - Microsoft Windows 7 (32-bit): windows7
      - Microsoft Windows 8 (32-bit/64-bit): windows7

- Microsoft Windows 7 (64-bit): windows7-64
- Linux (32-bit): rhel6
- Linux (64-bit): rhel6-64
- Other (32-bit): other
- Other (64-bit): other-64

For example, if the appliance is based on Microsoft Windows Server 2008 (32-bit), update the options field with the setting: esx_os_name=longhorn.

5. Save the application.
6. Start the appliance to verify that it is operational.
7. If the appliance was from a catalog, move the appliance back into the catalog.

   The appliance is now updated to use the new 3.x ADL descriptor format.

Update Appliance Boot Volume

Follow these steps to update an appliance's boot volume (Linux-based) to work on both Xen and ESX based servers within a 3.x grid.

1. Ensure your appliance is using the updated 3.x ADL descriptor format as specified in the previous section above.
2. Ensure the following statements are true about the appliance being converted:
   - The appliance has grub installed.
   - If grub uses /boot/grub/menu.lst for its configuration (this is the case with Ubuntu and Debian), then grub should not be updated after running this procedure as it will revert to the default grub configuration which will make the appliance unbootable. CentOS-based appliances are not affected.
3. Import the appliance into the 3.x grid using the class import command (or import your custom catalogs or applications as needed). Import anything that contains appliances that you want to update.
4. If the appliance is not a singleton, create a new application, drag an instance of your appliance into the application and branch the appliance. If the appliance is a singleton, edit the application in the infrastructure editor.
5. Create an empty partitioned boot volume. The size should be roughly set to the size of the existing boot volume (unless a different size is needed) + the size needed for a new kernel (usually about 50MB).
   - vol create my-app:boot_vol par1.size=150M
6. Copy the old boot volume's data to the partitioned volume that was created in the previous step.
   - vol copy my-app:LUX5.boot my-app:boot_vol%par1 --fscpy

7. Replace the old boot volume with the new partitioned volume.
   - vol copy my-app:boot_vol my-app:LUX5.boot --overwrite --force
   - vol destroy my-app:boot_vol -f
8. Manage the boot volume.
   - vol manage _GLOBAL_RO:apk_linux my-app:LUX5.boot -rw
9. Execute the following in the filer vol manage console to install a CentOS5 32-bit kernel (the same one that the Linux Filer uses). If you want a different kernel, copy it instead and update the grub configuration with the proper names. The boot volume is mounted under /mnt/vol2/par1.
   - cp -a /boot/{initrd-2.6.18-238.9.1.el5PAE.img,vmlinuz-2.6.18-238.9.1.el5PAE} /mnt/vol2/par1/boot/
   - cp -a /lib/modules/2.6.18-238.9.1.el5PAE /mnt/vol2/par1/lib/modules/
   - sed -i -e 's#/dev/hda2#/dev/hdb#' -e 's#/dev/hda3#/dev/hdc#' -e 's#/dev/hda4#/dev/hdd#' /mnt/vol2/par1/etc/fstab
   - tar -xvf /mnt/vol/apk-2.0.36-linux-rh.tar.gz -C /mnt/vol2/par1 (replace with the APK for your distro; verify you have the latest APK from the CA AppLogic® 3.x release)

- cd /mnt/vol2/par1
- tmp/apk-install
- sed -i -e 's/tty1/console/' /mnt/vol2/par1/etc/inittab

```
cat << EOF > /mnt/vol2/par1/boot/grub/grub.conf
default 0
timeout 1

title CA AppLogic® Appliance
      root    (hd0,0)
      kernel  /boot/vmlinuz-2.6.18-238.9.1.el5PAE root=/dev/hda1 ro console=tty0
console=ttyS0,38400n8
      initrd  /boot/initrd-2.6.18-238.9.1.el5PAE.img
EOF

fs=$(grep "^/dev/sdc1" /proc/mounts 2>/dev/null|awk '{print $3}')
[ "$fs" == "ext3" -o "$fs" == "ext2" ] && fs=e2fs

out=$(grub --batch --no-floppy 2>&1 <<EOF
device (hd0) /dev/sdc
root (hd0,0)
embed /boot/grub/${fs}_stage1_5 (hd0)
EOF
)

sectors=$(grep 'sectors are embedded' <<< "$out" |grep -Eo '[[:digit:]]+')

grub --batch --no-floppy <<EOF
device (hd0) /dev/sdc
root (hd0,0)
install --stage2=/mnt/vol2/par1/boot/grub/stage2 /boot/grub/stage1 (hd0) (hd0)1+${sectors} p
(hd0,0)/boot/grub/stage2 (hd0,0)/boot/grub/grub.conf
quit
EOF
```

Important! When vol managing the boot volume, update the fstab (/mnt/vol2/par1/etc/fstab) to mark any read-only volumes as read-only. If this is not done, the appliance will fail to boot.

10. Exit the filer console

11. In the infrastructure editor, edit the class and change the following:
    a. Change the device schema for your appliance to use /dev/hdX
    b. Update the device names for the volumes as follows:
        - /dev/hda1 -> /dev/hda
        - /dev/hda2 -> /dev/hdb
        - /dev/hda3 -> /dev/hdc
        - /dev/hda4 -> /dev/hdd
12. Save the application.
13. Start the appliance to verify that it is operational.
14. If the appliance was from a catalog, move the appliance back into the catalog.

    The appliance is now updated to boot on both Xen and ESX servers within a CA AppLogic® grid.

Note: The entire catalog and all applications that are included with CA AppLogic® 3.x have been updated to work on either Xen or ESX. You may choose to re-base your custom appliances on the CA AppLogic® 3.x appliances to obtain the ability to run them on either Xen or ESX. In this case there is no need to follow the instructions above.

Upgrading New Windows APK

If you participated in the beta, you can upgrade the new Windows APK (see page 15) from the version included with beta (APK version 3.5.4) to version included with GA (APK version 3.5.14).

To upgrade the Windows_APK.*.msi, remove the following directory links before installing the GA version:

- /lib/applogic
- /etc/sysconfig
- /var/run/applogic

What's Included

This section contains the following topics:

This release of the CA AppLogic® grid operating system includes the aforementioned key components. Each of them is described in a section that follows.

Distributed Kernel

The CA AppLogic® distributed kernel provides a set of system services required to support the distributed infrastructure and application model of CA AppLogic®. The four most important system services include:

- **Global volume store:** a scalable, distributed volume store using either the embedded virtual SAN on the grid servers or external remote SAN.

  If you are using the embedded virtual SAN, the volume store keeps volumes mirrored across two servers, ensuring high availability and improved read performance. If you are using external remote SAN, you must implement the mirroring on the external SAN.

- **Distributed virtual machine manager:** a run time component that virtualizes the hardware resources used by applications.

- **Logical connection manager:** a run time component that provides the virtual network bindings between components of an application without the need to configure any IP addresses and network settings for distributed applications

- **Application scheduler:** a run time component that selects and assigns hardware resources to applications, based on available grid resources, application constraints and user-provided configuration

Grid Dashboard

The grid dashboard provides the following features:

- **At-a-glance summary** of the grid state, including grid name, version, state summary, resource use, messages, settings, and so on.
- **List of currently installed applications**, with the ability to create new applications, copy existing applications, start and stop applications, and so on.
- **Log viewer** to view the grid controller logs, and also search, filter, and export functionality.
- **Support page** with links to user documentation, release notes, support forums, the Grid University, and so on.

Application Configurator

The application configurator is a control panel for configuring application parameters including setting their hardware resources, network resources, tuning, and other parameters. It is a single property sheet that includes all configurable parameters.

The application configurator can also be accessed through the command-line shell or scripts using the app configure command.

Infrastructure Editor

The infrastructure editor is a visual tool that makes it easy to create, assemble, and troubleshoot disposable infrastructure for CA AppLogic® applications.

The user interface of the editor is modeled after popular drawing programs where you assemble infrastructure by dragging components onto the canvas, connecting them together, and configuring each component using a property sheet.

For running applications, the editor can be used to open the monitoring dashboard for the application, and to start the grid shell for the application or login to individual appliances.

Command-Line Shell

The command-line shell gives you control of all aspects of a CA AppLogic® grid. The shell runs on the CA AppLogic® controller and can be accessed either through a browser, using the new web-based shell, or over SSH using any suitable SSH client package.

The shell commands are designed with the following objectives in mind:

- Make the shell easy to use
- Provide simple means for scripting automation

All commands have a batch form of their output that makes it easy to parse programmatically, while the command's default output is structured for convenient interactive operation.

Note: The CA AppLogic® application programming interface is also available. It provides a web service interface to one or multiple CA AppLogic® grids through a Representational State Transfer (REST) based service. The CA AppLogic® web srvices API enables developers of RESTful client software to directly interface with CA AppLogic® based Virtual Data Centers (VDC) (that is, a CA AppLogic® grid). The API allows for programmatic control of large pools of virtualized infrastructure to be available within each VDC. To use the Web Service API, the WS_API application should be running on a grid. The WS_API application provides HTTP, HTTPS and VPN based access to the API.

Application Programming Interface

The CA AppLogic® Application Programming Interface provides a web service interface to one or multiple CA AppLogic® grids through a Representational State Transfer (REST) based service. The API enables developers of RESTful client software to directly interface with CA AppLogic® based Virtual Data Centers (VDC) (that is, a CA AppLogic® grid). The API allows for programmatic control of large pools of virtualized infrastructure to be available within each VDC. To use the Web Service API, the WS_API application should be running on a grid. The WS_API application provides HTTP, HTTPS and VPN based access to the API.

Application Infrastructure Build System

The infrastructure build system compiles the application infrastructure, producing a single entity for the application. It verifies resource and configuration constraints for each appliance and for the application as a whole, builds instance images and enforces the integrity of the application infrastructure. The infrastructure linker binds the application instance to the grid hardware resources just in time for application start, producing a ready-to-run application from the portable application format.

The infrastructure build system is automatically invoked when starting applications and is transparent for the grid operator.

Application Monitoring System

The application monitoring system provides a visual interface for monitoring performance and resource usage statistics of running CA AppLogic® applications. The user interface of the Monitor is highly interactive and is accessible with a web browser.

System Catalog

The system catalog contains 30 appliance classes, ready to use in applications.

- **TOMCAT/TOMCAT64**: Tomcat application server (Sun Java machine and Apache Tomcat); 32-bit and 64-bit
- **JBOSS/JBOSS64**: Java EE-based application servers (32 and 64-bit)
- **ROR**: Ruby on Rails
- **WEB5/WEB64**: Apache-based web server with plug-in content/scripts volume
- **WEBx4**, **WEBx8**: Scalable web servers
- **MYSQL5**: MySQL-based database server
- **MYSQLR/MYSQLR64**: 32/64-bit MySQL-based database servers suitable for replication
- **PGSQL64**: PostgreSQL database server 64-bit appliance
- **ORACLE**: Oracle Express database server
- **NAS**: Network attached storage / file server appliance (HTTP and CIFS file access)
- **NASR**: Replicated network attached storage / file server appliance (HTTP and CIFS file access)
- **LOAD:** Load Generator that can be used to test various load scenarios in your CA AppLogic® applications
- **SQUID**: SQUID proxy (web cache)
- **HALB**: Session-aware HTTP load balancer based on HA Proxy
- **L3LB**: TCP/UDP load balancer based on HA Proxy
- **PS8**: Scalable port switch for distributing TCP and UDP traffic to different appliances
- **RPL**: Event replicator that replicates incoming HTTP requests to different appliances
- **URLSW**: URL port switch for distributing HTTP requests to different appliances based on a regular expression
- **MTA**: Mail transfer agent
- **INSSL**: HTTP Input Gateway with SSL Support

- **INSSLR**: Redundant HTTP Input Gateway with SSL Support (useful for disaster recovery purposes)
- **IN**, **OUT**, **NET**: Firewalled network gateways based on iptables
- **VPN**: Virtual Private networking appliance
- **MON**: Application Monitor used to monitor running applications (collects and displays counters using visual graphs)
- **LUX5/LUX64**, **LINUX5/LINUX64**: A tiny and a minimal Linux appliances that can be used as basis for new appliances

Use the following best practices to update your applications (if they are using any of the obsolete classes mentioned above):

- It is best to replace these appliances in your applications before the applications are migrated to your new 3.5 grid (assuming you are using 2.4/2.7/2.8/2.9/3.0/3.1 and these appliances exist on your grid). In this case, the appliances can be replaced by opening the application in the infrastructure editor, hold the SHIFT key and drag/drop the new appliance class over the existing appliance class on the canvas. The editor will prompt to verify you want to replace the class. Repeat this for all of the obsolete appliance classes in all of your applications. Doing it this way will preserve all property settings and connections within the applications. Afterwards, save your application and the updated application is ready to be used. Note that this will work for all of the obsolete appliances above except for HLB. For HLB, use the following recommendation.
- If you are migrating applications from an older grid such as CA AppLogic® 2.1 or do not have access to the newer classes mentioned above, you will have to update your applications after they are migrated to your new 3.5 grid. When you open your application in the editor, the editor will display a message stating that the appliance class is missing, and the appliance will disappear from the canvas. In this case, you will need to drag the new appliance class instances onto the canvas and re-parameterize/re-connect the appliances. Afterwards, save your application and the updated application is ready to be used.
- If you do not replace the obsolete classes within your applications, the applications will fail to start. If you open such an application in the infrastructure editor, the editor will display a message stating that the appliance class is missing. To resolve these issues, follow the recommendations mentioned in this section above.

Note: INSSL, the HTTP Input Gateway with SSL Support, is now assembled out of INSSLR instead of being a stand-alone appliance.

You can create windows-based appliances on your CA AppLogic® grid, but they are not included with CA AppLogic®.

- **WIN03S/WIN0364S/WIN08S/WIN0864S**: Windows 2003/2008 Server Standard Editions 32/64-bit
- **WIN03E/WIN0364E/WIN08E/WIN0864E**: Windows 2003/2008 Server Enterprise Editions 32/64-bit

- **WIN03DC/WIN0364DC/WIN08DC/WIN0864DC**: Windows 2003/2008 Server DataCenter Editions 32/64-bit
- **WIN03W/WIN08W**: Windows 2003/2008 Server Web Edition 32-bit
- **IIS03x/IIS08x**: Microsoft Internet Information servers (Standard/Enterprise/DataCenter/Web editions)
- **IIS03yx4/IIS08yx4**: Scalable Microsoft Internet Information servers (Standard/Enterprise/DataCenter/Web editions)
- **IIS03yx8/IIS08yx8**: Scalable Microsoft Internet Information servers (Standard/Enterprise/DataCenter/Web editions)
- **SQL08x**: Microsoft SQL Server database appliances (Web/Standard/Enterprise/Developer/Workgroup/Express editions)

The system catalog is a *global* catalog, containing appliance classes that can be used by all applications on the grid. You can see the full documentation for each appliance in the catalog reference. The system catalog is read-only for CA AppLogic® users and can be changed only by the grid maintainer.

CA AppLogic® also includes the following global catalogs:

- Dynamic: used for storing CA AppLogic® dynamic appliances, currently this catalog contains 3 classes:
  - **MIG**: enables the containing application to migrate or snapshot (non-live) itself to another grid
  - **BCK**: enables automatic application backup to external services
  - **SLA**: enables dynamic scaling of an application by starting and stopping other appliances within the application in accordance with a user-defined policy
- User: used for your own production-level appliances, freely modifiable by CA AppLogic® users, by default this catalog is empty

See the Appliance Catalog Reference Guide for a list of all appliances and their data sheets.

Notes:

- We removed the OpenSolaris appliances OSOL and OSOL64 (as well as the VDSes VDS_OSOL and VDS64_OSOL) from the catalog and no longer distribute them with CA AppLogic®. However, the OpenSolaris filer is still distributed with CA AppLogic®, and CA Technologies still supports these appliances and applications. CA Support can provide you access to the original OpenSolaris appliances and applications.
- OpenSolaris-based appliances only work on Xen-based grids currently and do not work on ESX-based grids.
- The Aldo ci and Aldo ai commands are no longer supported in CA AppLogic® 3.x. To import catalogs and applications into your grid (that is, system_ms that is shipped with CA AppLogic®), copy the catalog/application to your grid's impex volume and use the cat import and app import CA AppLogic® commands.

Sample Applications

This CA AppLogic® release includes 16 ready-to-use application templates.

The CA AppLogic® release includes the following Virtual Dedicated Server (VDS) application templates:

- Linux
  - **VDS_CentOS55**: 32-bit CentOS 5.5 VDS
  - **VDS64_CentOS55**: 64-bit CentOS 5.5 VDS

- Windows
  - **VDS_Win03S/VDS_Win08S**: 32-bit Windows 2003/2008 Server Standard Edition VDS
  - **VDS_Win0364S/VDS_Win0864S**: 64-bit Windows 2003/2008 Server Standard Edition VDS
  - **VDS_Win03E/VDS_Win08E**: 32-bit Windows 2003/2008 Server Enterprise Edition VDS
  - **VDS_Win0364E/VDS_Win0864E**: 64-bit Windows 2003/2008 Server Enterprise Edition VDS
  - **VDS_Win03DC/VDS_Win08DC**: 32-bit Windows 2003/2008 Server DataCenter Edition VDS
  - **VDS_Win0364DC/VDS_Win0864DC**: 64-bit Windows 2003/2008 Server DataCenter Edition VDS
  - **VDS_Win03W/VDS_Win08W**: 32-bit Windows 2003/2008 Server Web Edition VDS

  Notes:

  - You can create windows-based appliances on your CA AppLogic® grid, but they are not included with CA AppLogic®.
  - We removed the OpenSolaris appliances OSOL and OSOL64 (as well as the VDSes VDS_OSOL and VDS64_OSOL) from the catalog and no longer distribute them with CA AppLogic®. However, the OpenSolaris filer is still distributed with CA AppLogic®, and CA Technologies still supports these appliances and applications. CA Support can provide you access to the original OpenSolaris appliances and applications.
  - OpenSolaris-based appliances only work on Xen-based grids currently and do not work on ESX-based grids.

The CA AppLogic® release also includes the following preconfigured Linux-based infrastructure templates:

- **Lamp**: basic 2-tier non-scalable WEB application
- **LampX4**: scalable Lamp
- **LampCluster**: scalable Lamp cluster

The CA AppLogic® release also includes the following preconfigured Windows-based infrastructure templates (based on Windows 2003 Server):

- **WISA**: simple 2-tier non-scalable WEB application (Windows/IIS/SQL/ASP.NET)
- **WISAx4**: simple 2-tier scalable WEB application (Windows/IIS/SQL/ASP.NET)

Note: You can create windows-based appliances on your CA AppLogic® grid, but they are not included with CA AppLogic®.

The CA AppLogic® release also includes the following ready-made pre-installed application templates:

- **TWiki**: web-based collaboration platform
- **SugarCRM**: customer relationship management system
- **WS_API**: CA AppLogic® web service API

The applications are ready to run, requiring only network settings to be configured.

Note: The Aldo ci and Aldo ai commands are no longer supported in CA AppLogic® 3.x. To import catalogs and applications into your grid (that is, system_ms that is shipped with CA AppLogic®), copy the catalog/application to your grid's impex volume and use the cat import and app import CA AppLogic® commands.

Backbone Fabric Controller (BFC)

You use the Backbone Fabric Controller (BFC) to install and upgrade grids (instead of ALD). The BFC contains a simple-to-use web-based GUI application that you use to create and manage all of your CA AppLogic® grids within a single backbone. The BFC automatically downloads the latest CA AppLogic® releases and hotfixes.

Installation

This section describes how to download this release and install a grid.

This section contains the following topics:

Prerequisites

To install CA AppLogic®, you need a set of servers (1-128) connected with a gigabit Ethernet network and a designated Backbone Fabric Controller (BFC) server.

Important! Read at least the Backbone Fabric Controller (BFC) User Guide before choosing and setting up your servers and resources. Not reading or not following this document will likely result in a trial-and-error process which can be long and expensive. We want to verify your installation is successful from the first time. Contact CA Support if any of the requirements are unclear.

In addition, you will need an ssh keypair to be used to authenticate the grid maintainer. The public key must also be provided to CA, so that you can gain access to the CA download server. Please e-mail your *public* key to CA Technical Support.

Note: For more information about ssh keys, see the man page on ssh-keygen.

Downloading the Release

Contact your account manager for obtaining access to this release.

You can read the Backbone Fabric Controller (BFC) User Guide for information about how to install the BFC. This guide also explains how to download the latest CA AppLogic® release and hotfixes and how to install and manage all of your grids within a single backbone.

Notes:

- Verify you are running ssh agent with the key that you provided to CA for downloads. If you don't have a key or would like to use a different key, contact CA Support.
- For this release the default 'Download User' in the BFC should be **applogic**. You can make this change in the Administration page of the BFC. After installation, start the BFC, click Administration, then click the Versions tab. The user is specified in the Download User field.

Installing a Grid

The *Backbone Fabric Controller (BFC) User Guide* describes how to install a grid. Ensure that BFC is installed on your distribution server before installing a grid.

Note your Passphrase for Future Use

Important!

During the BFC installation, you are prompted to create a passphrase. This passphrase is required for upgrade and recovery. Because the passphrase is one-way encrypted it cannot be retrieved, so be sure to make a note of it for future use.

Product Characteristics

This section describes the latest product characteristics for CA AppLogic®.

This section contains the following topics:

Dimension Limits

CA AppLogic® has the following dimension limits

Key System Dimensions

- **Xen:** maximum of 128 servers per grid supported (tested on up to 30 servers)
- **ESX:** maximum of 80 servers per grid supported (tested on up to 30 servers)
- 31 grids per back-end LAN
- 1024 applications per grid, up to 1024 applications running simultaneously

Other Dimensional Limits

Note: For more information, see OS Support Limitations.

- Per application
  - 512 network interfaces per application
- Per appliance, common dimensions
  - 64GB RAM
  - 16 CPU (1600%)
  - 2000 Mbps bandwidth
  - 1 default network interface

  Note: Appliances may have multiple external interfaces.
- Per appliance: Linux
  - 32-bit and 64-bit
  - Volumes: Xen PV: 12, VMware ESX/Xen HVM: 4
  - External interfaces and terminals (including the default interface): Xen: 15, VMware ESX: 10

- Per appliance: OpenSolaris
  - 32-bit and 64-bit
  - Volumes: Xen PV: 12, VMware ESX/Xen HVM: 4
  - External interfaces and terminals (including the default interface): Xen: 8, VMware ESX: 10

    Note: OpenSolaris does not work on ESX-based CA AppLogic® 3.x grids
- Per appliance: Windows 2003/2008 Server (Standard/Enterprise/DataCenter/Web R2 Editions, Standard SP2 Edition)
  - 32-bit and 64-bit (R2 for 2008 server)
  - 4 volumes
  - External interfaces and terminals (including the default interface): Xen: 8, VMware ESX: 10
  - maximum CPUs and memory (Microsoft Windows limitations are also documented here: http://technet.microsoft.com/en-us/library/cc758523.aspx)
    - Web: 4 CPUs, 2GB memory
    - Standard: 4 CPUs, 4GB memory
    - Enterprise: 8 CPUs, 64GB memory
    - DataCenter: 16 CPUs, 64GB memory
- Per server
  - 1024 virtual volumes (counting each mirror as a separate virtual volume)
  - 255 shares (counting each mirror as a separate share)
  - 127 mounts (counting each mirror mounted as a separate mount; i.e., 64 if mirroring by two)
  - Xen: 50 appliances (CA AppLogic® internal limit)
  - ESX: 48 appliances (CA AppLogic® internal limit, 2 less than Xen to account for the CA AppLogic® virtual switch and I/O VM)
- Per virtual volume (volume size)
  - volumes used in Linux: up to 2TB-1MB
  - volumes used in Windows: up to 2TB-1MB
  - volumes used in Solaris: up to 1TB-1MB (1,048,575 MB, verified)

- Grid controller High-Availability
  - maximum of 1 primary server; where the grid controller VM is currently running
  - maximum of 7 secondary servers; where the grid controller may end up running in case of primary server failures
    - CA AppLogic® has been certified to use up to 7 secondary servers
- Network High-Availability (physical network topology)
  - maximum of 2 switches for each network; external and backbone (4 switches in total)
    - each pair of switches must be of the same make/model and they must be identically configured
    - each switch must have LLDP enabled (link layer discovery protocol)
  - maximum of 4 NICs per physical server (2 NICs for the external network, 2 NICs for the backbone network)
- Automatic Volume Repair
  - maximum of 1 volume repair scheduled per server

Note: If you need different dimensions, contact CA Support.

Hardware Compatibility

Adhere to the following guidelines:

- When selecting hardware to use for an ESX-based 3.x grid, the hardware must conform to all of the following HCLs:
  - VMware ESX HCL
  - BFC Control Node HCL
- BFC control node (distribution server)
  - Intel Pentium class processor, or compatible AMD processor (Athlon, Opteron), 2 GHz or better recommended (dual core or better)
  - 2 GB RAM
  - 100 GB hard disk space (to house multiple versions of CA AppLogic®)
  - 2 gigabit Ethernet adapters
  - Unless you plan to use the Bare Metal Installation feature, you must install CentOS 5.5 in advance.
  - YUM repository (external or local for OS updates)

- CA AppLogic® grid nodes
  - Cisco UCS
  - X86, multi-core (dual, quad, dual-dual, dual-quad), EMT-64, VT-enabled
    - ESX: minimum of 4 cores
    - Certified: Pentium P4, Intel Xeon/Xeon-Harpertown/Woodcrest/Clovertown/Core i3, Intel Nehalem (Core i7: Lynnfield 860), AMD Opteron, AMD Athlon64
    - Supported: Intel Pentium P4 or better; AMD Athlon or better
    - Intel hyperthreading is automatically disabled by CA AppLogic® and is not used
    - Requires either Intel VT or AMD-V for hardware virtualization support (HVM)
    - HVM is mandatory for running Solaris 10 and Microsoft Windows on your grid
    - HVM is mandatory for using iso2class to install new appliances based on different OS distros
  - 64-bit support (32-bit servers are not supported)
    - maintainers can only use 64-bit enabled servers in their CA AppLogic® grids
  - Memory
    - Xen: 2 GB Minimum, 4 GB recommended, 128 GB tested
    - VMware ESX: 4 GB Minimum, 8 GB recommended, 128 GB tested
    - Xen 4.1 Theoretical Maximum: 5 TB RAM per server
  - 80 GB IDE/SATA/SAS/SCSI/SSD HDD (250 GB or larger SATA drive recommended, multiple drives/server supported)

- HDD controllers
  - Certified:

    Intel Corp. 82801EB/ER (ICH5/ICH5R)

    Silicon Integrated Systems [SiS] 5513

    Advanced Micro Devices [AMD] AMD-8111 IDE

    nVidia Corporation CK804 IDE (rev a2)

    nVidia Corporation CK804 Serial ATA Controller (rev a3)

    Intel Corporation Ibex Peak 2 port SATA IDE Controller (rev 05)

    Intel Corporation Ibex Peak 4 port SATA IDE Controller (rev 05)

    Intel Corporation 82801JI (ICH10 Family) 2 port SATA IDE Controller

    Intel Corporation 82801JI (ICH10 Family) 4 port SATA IDE Controller

    Intel Corporation 631xESB/632xESB/3100 Chipset SATA IDE Controller (rev 09)

    LSI Logic / Symbios Logic SAS1068 PCI-X Fusion-MPT SAS (rev 01)

    LSI Logic / Symbios Logic SAS1068E PCI-Express Fusion-MPT SAS (rev 08)

    LSI Logic / Symbios Logic LSI1068E (rev 08)

    LSI Logic / Symbios Logic LSI2008 (rev 03)

    LSI Logic / Symbios Logic MegaRAID SAS

    LSI Logic / Symbios Logic LSI MegaSAS 9260 (rev 05)

    PERC H200 RAID Controller

    PERC H700 Integrated RAID Controller

    PERC H800 Controller Adapter

- Supported: most IDE, SATA and SCSI devices supported by CentOS 5.5 (excl. Adaptec AHA-15xx). Additional hardware requirements and compatibility details are available.

- CA AppLogic® has not been certified with any hardware RAID solutions and as such it is recommended to disable hardware RAID on all servers (each grid contains its own shared, mirrored storage pool; IP-based SAN). There are customers who are using CA AppLogic® with their own RAID solutions but these types of configurations have not been certified/tested by CA.
  - SAN support:
    - NFS only (Network File System Protocol) – no other protocols are supported at this time.
    - Ensure SAN is accessible from all servers within the same backbone on the 192.168.0.0/16 network, and that persistent routes are in place on the SAN so the servers can exchange network traffic with the SAN. The servers of a grid automatically configure themselves to access the SAN based upon the SAN configuration specified upon grid installation.
    - Ensure that SAN exposes a single NFS read-write share that is accessible from all servers within the same backbone on the 192.168.0.0/16 network. You can configure the NFS share to be *async* or *sync* (for ESX-based grids, only *sync* is supported).
    - Ensure that the NFS share has at least 500GB of free space for every grid that uses the share. For example, if the NFS share is to be used for five different grids, the share should have 2.5TB of free disk space.

- Two supported Gigabit or 10GB Ethernet adapters. For example, Intel E1000, BroadCom, AceNic, SysKonnect GE (one for backbone and one for external network).
  - For Network High Availability each node is required to have 4 NICs.

    Note: Network High Availability is supported when using 1G and 10G NICs.
  - Certified:

    Intel Corporation 82541GI/PI Gigabit Ethernet Controller

    Intel Corporation 82571EB Gigabit Ethernet Controller (rev 06)

    Intel Corporation 82574L Gigabit Network Connection

    Intel Corporation 82576 Gigabit Network Connection (rev 01)

    Intel Corporation 80003ES2LAN Gigabit Ethernet Controller (Copper) (rev 01)

    Intel Corporation 82572EI Gigabit Ethernet Controller

    Intel Corporation X520-T2 Dual Port 10GBase-T Copper Server Adapter

    Broadcom Corporation NetXtreme BCM5704 Gigabit Ethernet

    Broadcom Corporation NetXtreme BCM5715 Gigabit Ethernet (rev a3)

    Broadcom Corporation NetXtreme BCM5721 Gigabit Ethernet PCI Express (rev 11)

    Broadcom Corporation NetXtreme II BCM5709S Gigabit Ethernet (rev 20)

    Broadcom Corporation NetXtreme II BCM5716 Gigabit Ethernet (rev 20)

    Broadcom Corporation NetXtreme II 5709 Gigabit Ethernet

    Broadcom Corporation 5709C Gigabit Ethernet

    Broadcom Corporation NetXtreme II 5716 Gigabit Ethernet

    Realtek Semiconductor Co., Ltd. RTL-8139/8139C/8139C+ (as external 10/100 NICs)

    Neterion's X3100 Series 10GbE PCIe

    Emulex OneConnect 10Gb NIC (be3) (rev 01)

    Note: Ensure that Hotfix hf7516 is applied to the grid.
- Supported: most Gigabit Ethernet network adapters supported by CentOS 5.5. Additional hardware requirements and compatibility details are available.

- IPMI 1.5 support for server power control and mgmt strongly recommended (see below for details)
- Firewall access through port 22 (SSH) to grm.3tera.net and download.3tera.net for all servers

- To use VMware ESX in your CA AppLogic® grid, you must help ensure your hardware is supported by the ESX hypervisor (http://www.vmware.com/resources/compatibility/search.php) and by CentOS 5.5
- CA AppLogic® Grid Node BIOS settings: To allow grid servers to PXE boot on the BFC backbone network so that hardware discovery can take place:
  - System Security
    - AC power recovery = ON
  - Integrated devices
    - Backbone NIC set to PXE boot
    - Other NICs enable but not set to PXE
    - Boot Settings > Boot Sequence

      Backbone NIC must be before disk
  - Power controller setup screen (separately accessed during boot sequence, often with something like Ctrl-E)
    - IPMI over lan = ON

- Networking hardware: 2 non-blocking gigabit Ethernet layer 2 switches (one for external network, one for backbone network)
  - All ports must be on the same switch, no cascading
  - Must have as many ports as you have servers, plus one
  - The spanning tree protocol (STP) must be disabled on the switches
  - Power Management Network Switch (IPMI v1.5 or better) strongly recommended
  - If using a network HA configuration:
    - must have 2 of each type of switch (2 external switches and 2 backbone switches)
    - each switch must have LLDP enabled (link layer discovery protocol)

    Note: To use the VLAN tagging feature of CA AppLogic®, the switch must support VLAN's
  - Certified:
    - 10G switches

      Force10 S2410P - 24 port 10GigE XFP switch

      Cisco Catalyst 4900M

      IBM RackSwitch G8124E
    - 1G switches for network HA and non-network HA configurations

      Cisco WS-C3560E-24TD

      Cisco WS-C3560E-48TD

      Cisco WS-C3560E-24TD

      Cisco WS-C3750G-24TS-1U

      Cisco WS-C2960G-48TC-L

      Cisco WS-C2960G-24TC-L

      Cisco IOS Software; C3560 Software (C3560-IPSERVICESK9-M); Version 12.2(40)SE

- Power-control for servers (IPMI-based)
  - CA AppLogic® is certified to work with the following IPMI configurations:
    - 1. IPMI cards directly connected to the backbone configured to use non-routable IPs (10.x.x.x), separate backbone from the servers used in CA AppLogic® grids; accessible only through VPN on the internal network
    - 2. IPMI cards configured to use routable IPs; accessible from the external network using name/password (non-VPN)
  - The following IPMI configurations may work but have not been certified by CA:
    - Same as option #1 above except accessible through VPN on the external network
    - Same as option #2 above except that the IPMI is accessible from the internal network using name/password (non-VPN)

Note: If you have different network or storage devices, contact CA Support.

Software Compatibility

Adhere to the following guidelines:

- CA AppLogic® Hypervisor Support
  - Xen 4.1.2 (Linux kernel 3.2.2 stable)
  - VMware ESX 4.1 Update 2
- Appliances: all operating systems listed in this section were verified using iso2class (and hvm2pv when needed)

  Note: For more information, see OS Support Limitations.
  - Linux based
    - 32-bit/64-bit Linux OS (a mix of 32-bit/64-bit appliances may be running on the same grid)
    - Kernel 2.6.16-33 with Xen 3.0.4 support and all later kernels/Xen support
    - 32-bit Tested: CentOS 5.5/6.1, RHEL 5.3/6.1, Ubuntu 8.0.4/10.04.4/11.0.4, Debian 5.0
    - 64-bit Tested: CentOS 6.1, RHEL 5.3/6.1, Ubuntu 10.04.4/11.0.4, Debian 5.0
    - Supported: all Linux distros based on recent 2.6 kernel (see Linux distros verified by CA AppLogic® developers and users)
    - SUSE Enterprise APK is BETA; please report issues to CA Support

  - Windows 2003/2008 Server based
    - 32-bit and 64-bit
    - Windows Server 2003 R2, 32 and 64-bit: Standard/Enterprise/DataCenter/Web Editions
    - Windows Server 2008, 32-bit only: Standard/Enterprise/DataCenter/Web Editions
    - Windows Server 2008 R2, 64-bit: Standard/Enterprise/DataCenter/Web Editions
    - Windows 2003 Server 32/64-bit and Windows 2008 Server 32/64-bit are supported with the Halsign Turbogate PV drivers provided with the release
- Appliance volumes
  - File systems supported: ext2, ext3, ext3-snapshot, ext4, reiserfs, fat16, fat32, Solaris ufs, Solaris zfs, NTFS
  - Swap volumes are supported and optional for appliances
  - Integration services for other file systems are available
- CA AppLogic® GUI
  - Browser support: Microsoft Internet Explorer (on Windows), Mozilla Firefox (on Windows/Linux/MAC), Google Chrome (on Windows/Linux/MAC), Apple Safari (on MAC)
  - CA AppLogic® has been certified to work with the following browsers:
    - Internet Explorer 8 and 9
    - FireFox 11.0, 12.0, 13.0, and 14.0
    - Chrome 18.0
    - Safari 5.1

Note: If you are interested in other Linux distros, file systems, and software infrastructure, contact CA Support.

Chapter 3: Known Issues and Important Notes

This chapter contains information about known issues with CA AppLogic®. It also contains a section containing important notes about the installation, configuration, and use of CA AppLogic®.

This section contains the following topics:

Important Notes

1. ALD is no longer used to install/upgrade grids and to import catalogs and applications. Taking the place of ALD is the Backbone Fabric Controller (BFC). BFC is a simple-to-use web-based GUI application that is used to create and manage all of your CA AppLogic® grids within a single backbone. See the BFC documentation for how to download/install BFC and how to use it to manage your CA AppLogic® grids. To import catalogs and applications into your grid (i.e., system_ms that is shipped with CA AppLogic®), copy the catalog/application to your grid's impex volume and use the cat import and app import CA AppLogic® commands.

2. CA AppLogic® 3.x now supports the VMware ESX hypervisor, in addition to Xen. While CA AppLogic® 3.x maintains all of the features and functionality for both hypervisors, there are some important usage aspects that are specific to VMware ESX when using it with CA AppLogic®:

   - ESX has a more restricted hardware compatibility list as compared to Xen. Be sure to verify that your servers are on the ESX hardware compatibility list before using CA AppLogic® 3.x. The hardware choosen for your ESX-based grid must be verified against the HCLs for ESX, CA AppLogic®-Xen (specified in the hardware compatibility section above) and CentOS 5.5.
   - Heavily loaded grids that use ESX as compared to Xen operate slower for some operations. This is caused by the fact that CA AppLogic® uses the ESX APIs to control virtual machines on each server; these APIs are slower on ESX as compared to using the Xen APIs. The most noticable operations where this slowness is observed are the following:
     - Starting an application - appliances may take 1-2 minutes longer to boot on heavily loaded grids
     - grid info, srv list, srv info - may take several seconds or 1-2 minutes longer on heavily loaded grids
   - ESX has more memory overhead as compared to Xen. An application that fits on 2 servers on a Xen-based grid may not fit fully on the same 2 servers running on a CA AppLogic® ESX-based grid. This is why it is recommended for the ESX-based servers to have at least 4-8 GB of memory.

- Appliances that support ESX:
  - All appliances must have vmware-tools installed (by default, all appliances that ship with CA AppLogic® 3.x have vmware-tools installed). vmware-tools is needed so that the graphical console will work correctly and so that the appliance will shutdown in a timely manner. If vmware-tools is not installed, the graphical console is much harder to use (mouse cursor will be hard to control) and the appliance will take 15 minutes to shutdown.
  - When using Windows appliances, the appliance must have the correct esx_os_name setting in the virtualization options string that is stored in the appliance descriptor.You can update the setting (see page 25) for your windows appliance . If installing a new Windows appliance using iso2class, there is a mandatory command line parameter that is used to set the correct esx_os_name setting for the new appliance.
  - To reboot an appliance, use either comp restart or restart the application. Rebooting the appliance from within the appliance itself will result in an appliance boot failure (the appliance won't be able to retrieve its configuration from CA AppLogic®).
  - The volfix configuration mode will not work for appliances that run on ESX. The appliance will need to be converted to use dhcp configuration mode. This should only affect older appliances from CA AppLogic® 2.1 and older releases.You can convert your appliance to use the dhcp configuration mode .See the Appliance Developers Guide in the Appliance Kit section for how to convert your appliance to use the dhcp configuration mode.

3. CA AppLogic® 3.x introduces role-based access controls (RBAC). RBAC provides for the ability to grant permissions (control over) to an object (application template, application instance, catalog or grid). By default, when a new user is created on a grid, this user has limited access to the grid's objects. For example, by default the user does not have login permissions to the grid. You need to configure appropriate access rights for your users to access the grids.

4. When using the Web API to access your CA AppLogic® 3.x grid, the Web API application must be configured with a user who has full access rights to the grid (administrator-type access).

5. In CA AppLogic® 3.x, the operating system support for appliances is not as vast as it is for prior CA AppLogic® releases. In 3.x, Solaris 10 is not supported at all, and OpenSolaris is only supported for Xen. Note that OpenSolaris and Solaris 10 are being discontinued by Oracle and are being replaced by Solaris Express and Solaris 11.

   Note: All of the Solaris-based appliances have been removed starting from CA AppLogic® 3.5, excluding the Solaris filer. Contact CA Support if you need these appliances.

6. CA AppLogic® is OS agnostic and designed to be used with different operating systems. As part of its design, all volume operations (create/format, copy, resize, file-system check/repair and manage) are executed within CA AppLogic® applications named filers; they are no longer executed by the CA AppLogic® grid controller as with previous versions of CA AppLogic®. As such, these new filer applications use resources on the grid just like any other regular CA AppLogic® application. Therefore, there must be enough available resources on your grid to execute any of the CA AppLogic® volume operations. Note that the filer applications are not used for raw volumes or block-level volume copies.
7. The grid controller uses 10% of a core (for both ESX and Xen-based servers).
8. Because all volume operations are now executed using filer applications, all volume operations are slower as compared to previous CA AppLogic® releases since the filer applications have to be started/stopped as part of the volume operation. Typically, there is about 20 seconds of overhead for Linux-based volume operations and about 130 seconds for Solaris-based volume operations.
9. Network bandwidth resource usage is enforced on all appliances. An appliance will not be able to use more than its configured bandwidth for all of its terminals (the assigned bandwidth takes all terminals into account). Verify that the configured bandwidth for your appliances and applications is appropriate according to your bandwidth usage needs (or else you may experience some very slow network performance in your application). The maximum bandwidth per CA AppLogic® server is 2 Gbits.
10. Appliances that pass through network traffic such as gateways, load balancers and port switches, the bandwidth is actually cut in half. For example, a load balancer that is assigned 100M of bandwidth is actually limited to 50M (due to the fact that the network traffic passes in and out of the appliance).
11. Before accessing the CA AppLogic® GUI on a newly installed or upgraded CA AppLogic® grid, the user should clear out the browser's cache. If the browser's cache is not cleared out, the CA AppLogic® GUI may not behave properly.
12. The grid shell can be accessed either through a web browser or using an ssh client. For increased security, password-based ssh logins are not supported except during grid installation.

    Important! We strongly recommend that you use the web shell provided with the CA AppLogic® GUI.
13. When accessing the grid over ssh, the login user name is always root, regardless of the CA AppLogic® user name. For the purpose of ssh logins, users and their roles are uniquely identified by their public ssh keys.
14. Web browser's Javascript and pop-ups must be enabled to use the web-based graphical user interface (dashboard, editor, documentation)
15. Users are responsible for allocating, assigning and use of externally visible IP addresses for applications; CA AppLogic® takes care of all internal network assignments

16. While the Backbone Fabric Controller sets up all grid servers and controllers with carefully pre-configured firewalls and disables unnecessary network services, users and maintainers are encouraged to verify the security settings of their systems.
17. Network performance between servers on the private network used for volume and inter-appliance communication is measured to approximately 900 Mbps. The TCP network performance measured between appliances residing on different servers is measured as 720-900 Mbps. When running Windows, the TCP network performance is about 940 Mbps and UDP is about 500-700 Mbps.
18. Resource limits on appliance hardware resources are enforced differently for different types of resources (CPU, memory, bandwidth). CPU is "no less than" , memory is "exactly that much" (includes VM overhead), bandwidth is "exactly that much". CPU resources may be enforced to "exactly that much", using the new --cap_cpu option when starting the application.
19. When starting an application with a specified amount of minimum CPU, it is not guaranteed that the application will get exactly the amount of specified CPU. For example, if an application is started with cpu=2, it is possible that the application will receive 1.97 CPU as observed by adding up all of the assigned CPU to all components of the application. This is due to rounding errors that may occur while trying to assign CPU to each individual component.
20. When application start fails, not all messages related to the failure may be shown in the shell. Inspect the grid log for additional information, using the list log n=20 command.
21. Grids in which linear scalability of performance is important should be built using servers that are as uniform as possible in CPU type/speed, memory size and disk capacity. CA AppLogic® will work correctly in grids assembled from servers with different amounts of hardware resources; however, on such grids you may experience sub-linear performance.
22. There is no user visibility during a grid controller restart due to a grid controller VM failure. If the grid controller VM fails and CA AppLogic® restarts the grid controller VM, there is no user visibility while the controller is restarting. Typically it takes 1-2 minutes for the grid controller to restart on its own. If the grid controller is unavailable for more than 5 minutes, contact CA Support.
23. Creation of an NTFS03 volume always results in an NTFS08 volume. NTFS08 volumes may be used with Windows 2003 Server.
24. The net_discover command for grid and server is not supported on ESX-based grids/servers.
25. When using a SAN with your CA AppLogic® grids, ensure that there is at least 500GB of free space for every grid that uses the configured NFS share. For example, if the NFS share is to be used for five different grids, the share should have 2.5TB of free disk space.

26. When using a SAN with your CA AppLogic® grids, if the SAN or NFS share goes offline for any period of time, some of the volumes that were in use might get corrupted. If this corruption prevents the grid controller from running or causes applications to fail to start (or any other grid or application instability), contact CA Support immediately.
27. To use CA AppLogic® appliances based on the latest OS distributions (such as Fedora Core, Ubuntu, Debian, and CentOS), configure the new field engineering code of 128 on the boundary of the appliances. This new field engineering code instructs CA AppLogic® to use a newer device name style for the appliance volumes that are used specifically by these newer distributions. If the field engineering code of 128 is not specified, appliances based upon these newer distributions will fail to start.
28. If either your primary or replica BFC database is lost or corrupted, you may be able to recover it from an automatic backup that is always being run from the 3.1 version of the BFC. These backups actually live in a subdirectory of the primary database so they are **not** a substitute for configuring a replica. (These backups are also written to a subdirectory of the replica if you have one configured.) To restore from the most recent backup:
    - Wait two minutes after all active grid operations complete to ensure the backup is up-to-date.
    - Log in to the BFC system as root.
    - Run */bin/db_restore* (by default this will be */opt/bfc/bin/db_restore*). The db_restore utility will stop the BFC (if it is still running), restore the database from the most recent backup, and then restart the BFC.
    - After the restore you might find one or more of your grids in "Running, but needs attention" state. If you do, simply clear the failure on those grids before proceeding.

Known Issues and Limitations

This section describes the known issues and limitations at this time.

CA AppLogic® Limitations

1. Grid size is limited to 128 servers per grid for Xen, and 80 servers for ESX
   This is a limitation of the current CA AppLogic® release. This release has been certified up to 30 servers; however configurations up to 128 (for Xen) and 80 (for ESX) servers are supported.
2. CA AppLogic® currently cannot be installed on servers with 3 TB or larger drives.
3. For proper operation of the user interface, enable JavaScript, pop-ups, and cookies for the grid controller host. Ensure your browser is updated with the latest available version to correct defects that affect AJAX applications.

4. Protocols are not enforced on appliance terminals, only endpoints are enforced.

   This means that an appliance can only talk to appliances connected to it (plus its own server and the grid controller). Nevertheless, protocols on new appliances should be properly specified to help ensure application design integrity and compatibility with future versions of CA AppLogic®.

5. The total available disk space does not take volume mirroring into account

   The total available disk space reported by the grid info command is a raw estimate and does not take volume mirroring into account. The true available disk space is the reported available amount divided by the number of mirrors (2 mirrors by default). For example, if there is 1000GB of available disk space and the grid was configured for mirroring of 2, the available disk space is 500GB. Also, to successfully mirror volumes, there must be enough disk space on at least X servers where X is the number of mirrors (CA AppLogic® will not fail to create a volume if any one of its mirrors cannot be created, it will display a warning that the volume could not be mirrored).

6. A server failure during application start may cause the application start to fail

   If an application is started and one of the grid's servers fails, the application start will fail if one or more of the application's appliances were scheduled to run on the failed server. If this situation occurs, simply restart the application.

7. The volume management GUI available using the filers has an upload limitation of 10MB per file

   To upload larger files to your volume, use the vol manage shell command; don't forget to specify the external IP settings for this command to enable remote access from within the volume manager. For more information, see the reference for the vol manage command.

8. zFS may be used on the boot volume of an OpenSolaris appliance

   CA AppLogic® does support booting OpenSolaris appliances from a zfs-based boot volume. Please note however that this has not been verified by CA and may not work. Solaris 10 does not support zfs.

9. The CA AppLogic® Solaris filer that is used for Solaris volume operations supports a limited variation of zfs

   Currently this is limited to single device zfs pools. To take full advantage of all of the zfs capabilities in CA AppLogic®, users may assemble their own zfs pools inside of their own appliances. If a zfs pool is going to be used for mirroring, the CA AppLogic® volumes that are used in the pool should be created with the CA AppLogic® mirroring disabled (using the mirrored=0 option when creating the volumes). Also, a zfs pool created using the CA AppLogic® Solaris filer will not work in Solaris 10. See RefOsLimitations for all of CA AppLogic® OS limitations.

10. The maximum volume size for the ufssol file system is 1TB-1MB

    If you need larger storage, please use a different file system.

11. The property markup for appliance configuration is only supported for the volfix configuration mode

    The new dhcp configuration mode does not support the property markup for appliance configuration. When porting appliances from volfix to dhcp configuration modes, the APK documentation describes how to deal with appliances that depend upon the property markup for appliance configuration. See the Appliance Kit (APK) for more information.

12. Validation flags don't appear if the application is opened in read-only mode

    To see the validation flags for an application, open the application in edit mode. The validation flags are used to flag appliances that do not have all of their mandatory properties/terminals/volumes properly configured.

13. Text-based graphical console must be used with Solaris 10 after installation

    iso2class may be used to install a Solaris 10 appliance using the graphical console for the installation process. However, after the installation is complete and the appliance is re-started, the graphical console may still be used however it must be used in text mode (no access to the Solaris 10 desktop - strictly text-based access). This is due to a problem in the Solaris 10 GUI (not a CA AppLogic® bug).

14. All of the appliances that are distributed with CA AppLogic® are missing their GUI/desktop packages/support (X11, Gnome Desktop, etc.)

    Therefore, the graphical console cannot be used with these appliances. This is done on purpose to make the appliances as compact as possible. Using the new iso2class utility, users may create their own appliances with full desktop support.

15. Running more than 1 Windows application with appliances having the same instance name results in a duplicate computer name error from Windows

    This error is due to the fact that CA AppLogic® sets the computer name of an appliance to its instance name. Therefore, if you have more than 1 appliance running on a grid that all have the same instance names, the duplicate name error will be displayed in Windows on the graphical console. This error is simply a warning and does not affect the grid or its operation. However, if you need to use Windows as a domain controller, you will need to set the computer names to unique names for each appliance. You may use the wincfg utility to set the computer name in your appliance.

16. The graphical console requires the latest version of Java in your IE/FF browser

    We have tested with Java version 6 update 7 on IE/FF/Chrome/Safari. If the latest version of Java is not used, the graphical console may not work correctly (it will hang while trying to load). Before reporting graphical console errors to CA, be sure to verify that you are using the latest Java version (if you need to upgrade java in your browser, be sure to re-open your browser afterwards for the graphical console to work correctly).

17. Failover groups may not be satisfied upon controller recovery

    When a secondary server takes over as the new primary server, if there are not enough resources available on the server to start the grid controller, CA AppLogic® restarts appliances which are running on the new primary server on other servers within the grid so the grid controller can be started on the new primary server. Note that this may break appliance failover groups. If CA AppLogic® stops one of these appliances it may not be able to restart the appliance on another server because there may not be enough resources to satisfy the failover group.

18. Xen-HVM-based appliances use more memory than their configured amount

    All HVM-based appliances (Solaris 10, Windows, etc.) use more memory on the server than what they are configured to use. Typically, depending upon the amount of memory assigned to an HVM-based appliance, the appliance uses additional memory on the server in which it is running (this additional memory is required by the virtualization hypervisor running on the servers and is known as shadow memory). Therefore it is possible that even though a server might have enough available memory as compared to what is assigned for the appliance, the appliance will not be able to run on that server due to the additional shadow memory needed for HVM-based appliances that is not available on the server. The CA AppLogic® scheduler does take this extra shadow memory into account when scheduling appliances during application start.

19. The maximum network throughput between appliances running on different servers can range from 2 Gbps when using a 10G backbone (Xen and ESX).

    When using a 10G backbone, the maximum throughput that can be achieved between appliances running on different servers is about 2Gbps (possibly due to some sort of limitation within the hypervisor used by CA AppLogic®).

20. Internet Explorer cannot be used to access the graphical console of an appliance that is based on Ubuntu

    Any other browser may be used instead.

21. The new shared interface support in CA AppLogic® 3.x will not work on Windows-based appliances

    Shared interfaces should work with all other operating systems.

22. If any server of a grid has greater than 4 TB of local storage, the dom0 memory for the entire grid must be set to 1 GB. This is specified through the BFC using the parameter *dom0_vm_mb=1024* when you install the grid.

23. While the BFC lets you configure multiple IP ranges per VLAN, the CA AppLogic® Dashboard and Application Connections Configuration GUI lists every possible combination of VLAN and IP range. This will be fixed in a future release.
24. The catalog and sample applications that are distributed with CA AppLogic® have not been updated to take the new multiple external interface (MEI) feature into account. Therefore, if a user branches one of the default catalog appliances and wants to use more than one external interface, the appliance's scripts must be changed (the scripts assume that the appliance only has one external interface). The catalog and sample applications will be updated to support MEI in a future release.

Known Problems and Issues

The following are the known issues in this release:

Severity 1

- None.

Severity 2

- None.

Severity 3

1. Reduced Performance When Using HP Smart Array RAID Controller Without Write Cache

   When using the HP Smart Array RAID controller without the write cache enabled, there is a 50% reduction in performance. This issue has been verified on a HP DL 580 G7 Server, with Smart Array P410i 256mb. These cards require a battery or capacitor to be installed to enable the write cache.

2. Ensure the SR-IOV BIOS setting on Emulex 10G NICs is DISABLED

   When using ServerEngines Corp. Emulex OneConnect 10Gb NIC (be3) (rev 01) NICs with CA AppLogic®, these NICs incorrectly bounce packets if the SR-IOV BIOS option is enabled. These bounced packets alter the bridge's forwarding cache, causing the bridge to drop packets instead of forwarding them to the correct destination. This causes instability in CA AppLogic® which results in intermittent application start failures. Therefore, please ensure that the SR-IOV BIOS setting is DISABLED for all Emulex 10G NICs on all servers within the grid.

3. Defect SCR 6280 - Maximum network throughput between appliances running on different servers is approximately 2.5 Gbps when using a 10G backbone (Xen)

   When using a 10G backbone, the maximum throughput that can be achieved between appliances running on different servers is approximately 2.5 Gbps (you may observe different results depending upon the type of 10G hardware that is being used). CA is currently researching several network optimizations (such as enabling jumbo frames) that may be enabled in future CA AppLogic® releases in order to enhance 10G network performance.

4. Defect SCR 2203 - Stuck volume mount causes failures to start applications (rare, cannot reproduce)

   Very rarely an application will fail to start due to a stuck volume mount on one of the servers. CA AppLogic® detects stuck volume mounts and reports them to the user on the grid's dashboard. If this problem occurs on your grid, notify CA Support. Optionally, disabling the server or rebooting the server that has the stuck mounts will resolve this issue.

5. Defect SCR 3416 - Grid controller does not automatically recover if the primary server becomes unresponsive

   If this situation occurs, rebooting the primary server will restore the grid to an operational state.

6. Defect SCR 2239 - CAT

   VDS: security vulnerability: initial user/password setup

7. Defect SCR 2243 - The CA AppLogic® GUI leaks memory when used with Microsoft Internet Explorer 6 or 7

   If the CA AppLogic® GUI is accessed using Microsoft Internet Explorer 6 or 7, the GUI leaks memory as applications are opened for editing or when the web shell is opened (5-20MB of system memory are leaked for each of these operations). It is recommended to close and re-open the browser every few hours to recover the leaked memory. Firefox, Chrome or Safari may also be used instead of Internet Explorer.

8. Defect SCR 1471 - GUI times out and logs out the user while there is load on the grid controller

   The GUI no longer automatically logs the user out when there is heavy load on the grid controller. Instead, the user will receive a message stating that there was a network error. In this case however, the GUI is still fully functional. The network error message will only be received when there is heavy load on the controller, such as starting 4 applications at the same time AND copying a large multi-GB volume. In large grids, try assigning up to a full CPU core and 1GB RAM to the controller.

9. Defect SCR 857 - Grid reboot may degrade one or more system volumes

   If a grid is rebooted using the grid reboot command, when the grid comes back up after the reboot, one or more of the system volumes may become degraded. CA AppLogic® automatically repairs these volumes as highest priority.

10. Defect SCR 1199 - Unable to migrate a volume whose streams are all on disabled servers

    When migrating a volume, verify that at least one of its streams is on an enabled server or else the migration command will fail. The volume can be completely migrated off of its original set of servers by migrating the volume twice.

11. Defect SCR 1496 - Grid automatic application recovery (HA) may fail due to servers taking too long to reboot

    Some physical servers may take a long time to reboot - this may cause CA AppLogic®'s automated grid recovery to fail. The end result of this is that applications may not be all restarted automatically after the grid recovers from a failure. This is due to the grid controller waiting for a maximum of 10 minutes for all servers to reboot and reconnect to the grid controller (which may not be enough time for all servers to reboot). Workaround is to manually restart applications after all servers have reconnected to the grid controller - execute "list srv" to help ensure that all servers are connected to the grid controller - they all should be in the UP state. In CA AppLogic® 2.1, with server boot timeout of 10 minutes, this may occur primarily if a server fails to boot due to hardware or BIOS malfunction.

12. Defect SCR 1234 - Grid flapping file is not always reset when the operator intentionally reboots the grid

    When the operator reboots the grid, the grid flapping state is supposed to be reset and a message should be displayed on the dashboard stating that the operator rebooted the grid intentionally ("Grid has been restarted by operator on ..."). Occasionally when rebooting the grid, the grid file is not reset nor is the dashboard message displayed. The only problem that this may cause is upon the next grid failure, the applications may not be automatically restarted (depending on how many times the grid has failed when this bug occurs). To workaround this problem, if after an intentional grid reboot there is no dashboard message displayed, contact CA Support to have the grid flapping state reset on your grid.

13. Defect SCR 1360 - Appliance shows slightly less memory and less disk size than allocated

    The reason for the slightly reduced resources is related to allocation for service areas. For memory, it is likely due to Xen related to the memory map table for a virtual machine. For disk, it is due to normal file system service areas (this is the same as on regular Linux servers).

14. Defect SCR 2293 - Occasionally opening an application in the editor results in a message that the application is locked for editing

    In this case, the application is not opened for editing by any other user but the CA AppLogic® editor erroneously thinks somebody else has the application open for editing. If this occurs, simply override the application lock when prompted by the editor upon opening the application.

15. Defect SCR 2313 - IE6/7 are about 2x slower than FireFox/Chrome/Safari when using the CA AppLogic® GUI

    The main slowdown occurs when opening an application in the CA AppLogic® infrastructure editor.

16. Defect SCR 2497 - It takes 15 minutes to re-open graphical console after client computer crashed while graphical console was open

    If the client has the graphical console open and they lose connection to the internet (client network card failure, client computer crash, internet access is unavailable, etc.), it will take 15 minutes to re-open the graphical console.

17. Defect SCR 2548/ SCR 2549 - Issues when using the CA AppLogic® graphical console with Ubuntu

    The mouse is hard to use in Ubuntu when using the CA AppLogic® graphical console. This is due to a limitation of the Xen VNC support (mouse acceleration is not supported). Some users report that adjusting the mouse settings in Ubuntu resolves the issue. Also, rarely keystrokes will be repeated several times when typing in text from the keyboard (in such cases, simply delete the extra characters that are displayed).

18. Defect SCR 2498 - All text entered by the user in the text boot console is echoed to the console

    This includes passwords when logging into an appliance. The text boot console should only be used for debugging purposes. The SSH console can be used instead for all other purposes.

19. Defect SCR 2501 - User must press enter to see output in the text boot console after it is opened for the second time

    If a user re-opens the text boot console for an appliance after it has already been opened, they must press the enter key to see either the login prompt or the command prompt. This is because the boot console is waiting for user input (either for login information or a command to be executed).

20. Defect SCR 3107 - Appliances in failover groups are not accounted for when restarting the grid controller on a secondary server

    If a grid has an appliance that is part of a failover group running on a secondary server where the grid controller needs to be restarted, CA AppLogic® may stop that appliance which could break the failover group.

21. Defect SCR 2134 - Grid upgrade causes an incorrect warning about the cause of the grid reboot

    After upgrading a grid to the latest release, a dashboard message is posted stating that the grid failed due to a hardware issue. This message can be safely ignored and removed from the dashboard.

22. Defect SCR 3499 - The CA AppLogic® APK does not work with latest Ubuntu 9.10 or 10.x releases

    The appliance kit (APK) does not currently work with Ubuntu 9.10 or 10.x due to several incompatibilities with the newer OS. However, there are various posts on the CA AppLogic® forums that describe how to use some of the later OS distributions with CA AppLogic®.

23. Defect SCR 3709 - Appliances become temporarily inaccessible (5 min) if external NIC fails in network HA configuration

    If using a network HA configuration with CA AppLogic® and there is an external network failure, applications/appliances that use external interfaces may become inaccessible for up to 5 minutes. This appears to be caused by the external router caching MAC addresses. Waiting for the router to flush its ARP cache or sending an ARP response with arping from the application restores operation. This only affects the external network (the backbone network is not affected).

24. Defect SCR 4028/ SCR 4030 - Solaris 10 is not supported

    Solaris 10 does not work on CA AppLogic® 3.x for both Xen and ESX servers.

25. Defect SCR 4053 - OpenSolaris is not supported on ESX-based servers

    OpenSolaris only works on Xen-based servers.

26. Defect SCR 4159 - Recovery GUI does not work on ESX-based servers

    The recovery GUI only works on Xen-based servers.

27. Defect SCR 4273 - MON counters for shared interfaces do not work properly

    Shared interfaces do not support appliance counters.

28. Defect SCR 5242 - System Up time is not reset after the Grid power cycle is done

    If a user power-cycles a grid, the system uptime is not reset. If the grid is rebooted, the system uptime should be reset.

29. Defect SCR 5269 - The grid power_cycle command may fail to reboot the primary server

    If a user power-cycles a grid using the grid power_cycle command, the primary server may fail to reboot. This only occurs when the command is executed after a new grid install and the grid was never rebooted before the power cycle command was executed. Rebooting the grid at some point after a new grid install will avoid this issue.

30. Defect SCR 6592 – SAN-NFS: CA AppLogic® should automatically detect changes in the NFS share size during runtime.

    If the NFS share size is changed while a grid is running, CA AppLogic® will not detect this until the grid is rebooted. This issue will be resolved in a future release.

31. Defect SCR 6378 - SAN-NFS: Deletion of running/failed grid leaves behind the grid's folder (but contents of the folder are deleted)

    When a grid that used a SAN is destroyed, CA AppLogic® deletes the contents of the grid’s folder on the SAN, but leaves behind the empty folder. This issue will be resolved in a future release.

32. Defect SCR 6448 - Add server to Xen Grid fails with a kernel crash (Dell H200 RAID cards)

    Dell-based servers that use the H200 RAID cards cannot be used with CA AppLogic®. This issue will be resolved in a future release.

    The workaround for this problem is to enable hardware RAID on the Dell server before using it for grid creation.

33. Defect SCR 6684 - RedHat 5.3: appliance creation using iso2class utility with 4 iso images fails

    RedHat 5.3 based appliances cannot be installed using iso2class. This issue will be resolved in a future release.

34. Defect SCR 6701 - Upgrade from 3.0 or 3.1 to 3.5 rarely fails

    Very rarely, an upgrade to 3.5 from either 3.0 or 3.1 may fail. In this particular upgrade failure case, the following messages are present in the grid's status log accessed using the BFC (click on the status of the grid to open the log).

    installing the controller image
    ioctl: LOOP_SET_FD: Device or resource busy
    installing new controller FAILED, aborting

    If these messages are present in the log, rerun the upgrade again and it should succeed.

    Note: This issue is actually a bug in both CA AppLogic® 3.0 and 3.1, and is resolved in CA AppLogic® 3.5.

35. Defect SCR 7048 - ESX grid rollback from 3.5 to 3.1 does not work

    The rollback command does not work from 3.5 to 3.1 for an ESX-based grid. However, as a workaround, the downgrade command can be used (note that downgrade takes a bit longer than rollback). This issue will be resolved in a future release.

36. Defect SCR 7064 - ext3-snapshot based volumes do not work on ESX-based grids

    Ext3-snapshot based volumes do not work on ESX-based grids. However these volumes work on Xen-based grids. If you are using an ESX-based grid and you need to use an ext3-snapshot volume, you can add a Xen-based node to your grid and use that node to create/manage your ext3-snapshot volumes (when running the volume commands, disable all of the ESX servers so the CA AppLogic® filer will run on the Xen-based node). This issue will be resolved in a future release.

37. Defect SCR 7397 - "vol migrate" failed to migrate one of the volume streams from the local SAN (incorrectly trying to migrate stream to external SAN)

    An attempt to migrate a volume stream on the local SAN might fail on grids that are configured to use an external SAN. Iinstead of migrating the volume stream to the local SAN, CA AppLogic® incorrectly tries to migrate the stream to the external SAN. If you encounter this failure, use the **store=local** option with the **vol migrate** command. This issue will be resolved in a future release.

38. Defect SCR 7394: Grid controller hang after upgrading CA AppLogic® from 3.0.30 to 3.5.x

    When CA AppLogic® is upgraded from 3.0.30 to 3.5.x, the grid controller intermittently hangs and any 3tshell command executed returns a **low memory condition** error message.

    To work around the issue, reboot the grid controller. This issue will be resolved in a future release.

39. Defect SCR 6779: CA AppLogic® installation fails on some type of Broadcom nics

    On some Broadcom nics, and particularly NetXterme II BCM5709/5716, the link speed is reported as 100Mb/s or 10Mb/s by the nic driver. As a result, CA AppLogic® installation fails.

    To work around the issue, attempt a reinstall. This issue will be resolved in a future release.

40. Defect SCR 6517: Limit of 10 concurrent asynchronous jobs in the web service API.

    The OpenSSH version installed on the grid controller limits the number of simultaneous multiplexed ssh sessions to 10. As a result, if more than 10 asynchronous requests are executed, they are dropped by the API.

    To workaround the issue, issue less than 10 simultaneous asynchronous requests to the API. This issue will be fixed in a future release.

41. Defect SCR 7059: Infrastructure editor fails to load after renaming the Assembly or Component interfaces when using multiple external interfaces.

    If you rename an Assembly or component interface, the infrastructure editor does not load completely. This issue will be fixed in a future release.

42. Defect SCR 7516: Nic is reported as ‘active-down’ for ServerEngines Corp. Emulex OneConnect 10Gb Nic (be3) (rev 01) .

    On servers with these nics, after a grid is created, the output of **srv info srvX –extended** shows the state of the nics as **active-down**. This has been identified as a hardware specific issue. In order to work around the issue, log into the respective switch, shut the port of the nic on the srvX and enabled it again. The state should now be shown as up. This issue will be resolved in a future release.

43. Defect SCR 7507: Grid installation failures when using Dell R710 servers with Broadcom NetXtreme II 57711 Nics

    It has been observed that on Dell PE R710 servers with Broadcom NetXtreme II 57711 (bnx2x) 10Gbe nics, BFC fails to discover the servers leading to a failed install. This is a hardware specific issue and will be resolved in a future release.

Known Problems and Issues Specific to Windows-based Appliances

The following are the key known problems with Windows appliances in this release. Also, see the Windows Appliance Installation Reference for additional procedures and notes.

1. Defect SCR 6515 - APK (new-native): WIN0364DC appliance is intermittently failing to start on a Xen grid

   When using the new Windows APK that ships with CA AppLogic® 3.5, Windows 2003 Server 64-bit Data Center edition may intermittently fail to start when used on a Xen-based grid. If this issue is encountered, restarting the appliance may work around the issue. This issue will be resolved in a future release.

2. Defect SCR 2751 - Windows filer volume resize can fail on a volume with a corrupt file-system

   The Windows filer can fail a volume resize operation if the source volume contains a corrupt directory entry/file. The main source of this problem comes from the fact that some of the Microsoft software installations purposely contain invalid directory entries (we are not sure why this is; this has been observed when a user installed a version of Microsoft SQL Server in their appliance). Additionally, the source volume can be corrupt due to normal wear and tear. This issue can be worked around by running a file system repair on the volume (vol fsrepair) before resizing the volume.

3. Defect SCR 3078 - Resizing an NTFS volume failed due to a Windows filer start failure

   It has been observed by CA that the NTFS volume resize operation fails about 2 times out of 100. These 2 failures occurred because the Windows filer failed to start correctly on the grid. If this issue is observed, repeating the resize operation a second time should succeed. This issue however should be resolved in this release; if this issue is observed, notify CA technical support.

4. Defect SCR 2750 - Windows filer failed to create an ntfs volume (rare diskpart error)

   The Windows filer uses a Microsoft utility named diskpart to deal with the Windows NTFS volumes. Occasionally diskpart fails to obtain volume information or may fail to mount the volume. This is a very rare failure and may cause either vol create or vol resize to fail over NTFS volumes.

5. Defect SCR 2748 - Windows appliances occasionally detect duplicate IPs on their internal network

   If the user has an application that contains a Windows appliance and one or more Windows appliances are added to the app or terminals are added/removed from the Windows appliances, during the first app start some of the Windows appliances may detect duplicate IPs on their internal network (this can only happen during the first app start after the application is modified). This should not cause any operational failure of the application or require user intervention; the duplicate IP addresses are purely temporary. Worse case, some of the network communication involving any of the Windows appliances may be delayed for up to 30-60 seconds.

6. Defect SCR 3021 - Windows application stuck at 99% during application stop

   An attempt to stop a Windows application hung at 99%; the operation timed out after 15 minutes. The application contained 2 instances of a Windows 2003 Server DataCenter Edition appliance (WIN03DC). One of the Windows appliances stopped and the other one hung during "comp stop". This was only observed one time and could not be reproduced.

7. Defect SCR 2504 - Occasionally disk read/write counter values are reported as zero (Windows perfmon API bug)

   Occasionally zeros are reported for the following disk I/O counters for Windows appliances (even though sustained I/O is being generated): Total bytes written/read, # of volume writes/reads, time spent in writes/reads. This is due to a bug in the Windows perfmon API - the zero values is what is being reported by the Windows perfmon API.

8. Defect SCR 2821 - the Windows filer MSI do not work under localized Japanese Windows

   Other than the filer MSI, localized Japanese Windows should work under CA AppLogic®.

9. Defect SCR 2862 - Windows appliance fails to start if a virtual DVD-ROM device is installed

   A windows appliance fails to start if the MagicISO virtual DVD-ROM device is installed. Virtual DVD-ROM devices are not currently supported in CA AppLogic® for windows-based appliances.

10. Defect SCR 2499 - It can take several minutes to discover new NICs in Windows appliances which can cause boot timeouts

    Occasionally it takes several minutes for Windows to detect new NICs inside of an appliance. This occurs when the user adds/removes terminals for a Windows appliance singleton. The extra time it takes to detect these new NICs may cause appliance boot timeouts. To workaround this, increase the boot timeout of your Windows appliance.

11. Defect SCR 2505 - Migration of a windows appliance to another grid may trigger re-activation of the Windows appliance

    If a user has a Windows appliance on their grid and they migrate the appliance to another grid that has different hardware, the Windows appliance may require re-activation (Microsoft's Windows re-activation). The re-activation is triggered when a specific amount of hardware has changed (it is unknown to CA exactly what hardware changes trigger the re-activation). Note that re-activation may require access to the internet from within the Windows appliance. This particular problem was observed after resizing the Windows appliance boot volume and migrating the appliance to a different grid.

12. Defect SCR 3814 - Windows 2008 filer root access permissions are limited using ssh

    This issue only affects Windows 2008 Server 32/64-bit (Windows 2003 server works OK). When accessing a Windows 2008 volume either through the filer using ssh to an appliance, the user may not be able to access/modify files due to permission issues. To access/modify files using the command shell, log in through the graphical console to the Windows desktop and open up a command shell. The command shell can be used to access/modify files.

13. Defect SCR 4593 - Windows 2003 VDS/base classes fail to start (timeout on boot)

    Windows 2003 Server times out during its first boot during installation. Be sure to follow the windows build instructions to workaround this issue.

14. Windows 2003 Server based appliances - installing the Turbogate PV drivers requires user intervention

    When installing the Turbogate PV drivers, upon first start of the appliance when running on a Xen-based grid server, the user has to manually click through the hardware setup wizard for the installation of the Turbogate PV drivers for all terminals configured in the appliance. Otherwise, the appliance will fail to start.

15. Defect SCR 5737 - Windows 2003 Server 32/64-bit based appliances; only works on the hypervisor where the appliance was initially created

    When creating a new 32/64-bit Windows 2003 server appliance, the appliance will only work on a grid server that uses the same hypervisor on which the appliance was initially created. Otherwise, the appliance crashes during boot. For example, if the appliance is initially created on an ESX-based grid server, the appliance can only be used on an ESX-based grid server (attempting to use the appliance on a Xen-based grid server will not work, the appliance will crash during boot).

16. Defect SCR 5960 - Counters for a Windows 2003 appliance do not show up when using MON

    This is a known issue with Microsoft Windows 2003 server. Microsoft has a solution to resolve this issue with your Windows 2003 appliance.

Unreproducible Issues

The following issues have been observed in CA AppLogic® releases but are extremely difficult to reproduce (if at all) and have only been observed once or twice. If any of these issues appear on your grid, please send a bug report to CA describing which problem occurred and which CA AppLogic® commands were executed that led to the failure.

1. Defect SCR 2842 - Server rebooted due to a crash in the Linux kernel (observed in various releases)

   A server in the grid rebooted on its own due to a crash in the Linux kernel in dom0 of the server. This would not cause the entire grid to fail like in previous CA AppLogic® releases; but could cause application downtime. In such a case CA AppLogic® restarts the appliances that were running on the failed server on other servers in the grid. If this issue is observed on your grid, contact CA Support.

2. Defect SCR 2834 - Server loses connection to the grid controller

   In CA AppLogic® 2.4, there have been several cases where a server loses connection to the grid controller and reboots. This causes all of the appliances that were running on the server to be rescheduled on other servers in the grid and can also cause application downtime. It is unknown why the servers are losing their connections to the grid controller.

   If the server's connection to the grid controller is dissolved, the server tries to reconnect to the grid controller and if successful, the server remains operational and there is no application downtime. If the server cannot reconnect to the grid controller for one minute, the server is rebooted and application downtime occurs. When a server loses its connection to the grid controller, a message is logged to the dashboard. If this problem is observed, contact CA Support.

3. SCR 2903 - Volume resize of 4 NTFS volumes executed at the same time failed

   On CA AppLogic®, resizing 4 NTFS volumes at the same time caused all four volume resize operations to fail. This issue has been observed only once.

4. SCR 3289 - NASR replication failure was observed when almost out of disk space

   While NASR was replicating an 800MB file on a 1GB volume, the NASR appliance became unresponsive. CA is unable to reproduce this issue. If this issue is encountered on your grid, contact CA support.

5. SCR 3711 - Opening many graphical consoles crashed a server in the grid

   User opened six graphical consoles to different windows appliances running on the grid (opened at the same time). Upon opening the seventh graphical console, one of the servers rebooted and rejoined the grid. The appliances that were running on the failed server were restarted on other servers within the grid. This issue has only been observed once.

BFC Known Problems

We have identified the following known problems with the Backbone Fabric Controller (BFC) in this release:

1. If you are running a BFC database replica on an NFS hard mounted file system (NFS hard mounts are the default; do not use the optional soft mount functionality), and that NFS-mounted file system fails, the BFC will hang. This issue is a characteristic of NFS itself, and not something the BFC has direct control over. If you end-up in this state and you are unable to restore the NFS file system, you can remove the BFC dependence on that replica to restore normal operation using the following steps:

   a. Log in to the BFC system as root.

   b. Change to the bfcadmin user by typing the following command:

      su - bfcadmin

   c. Run *<BFC install location>/bin/stop_replication* (by default, */opt/bfc/bin/stop_replication*)

   Important! After breaking this dependence, your system will be running without replica, so go back into the UI and establish another replica at the same or a different location.

2. Defect SCR 6990 - Cannot unset the default VLAN for a grid via the BFC API

3. Defect SCR 6027 - The grid start from BFC UI fails after it is shutdown using “3t grid shutdown” command

   Please do not use the “3t grid shutdown” command on a grid.

4. Defect SCR 7036 - ESX grid fails due to nfs mount error

   When this occurs, doing a “service nfs restart” on the BFC should resolve the problem.

5. Defect SCR 7058 - Failed ESX grid node goes into an infinite loop of reboot once it is restarted

6. Defect SCR 6424 -BMI Install prompts for Driver Disk on HP DL360g4p

   If you get this message, simply hit the “Esc” key to continue the install.

7. Defect SCR 6779 - Servers that are known to have GigE interfaces sometimes report/fail saying that they are not running at GigE speeds

   In CA AppLogic® 3.5, some Broadcom Corporation NetXtreme II NICs misreport as being too slow. If you get this error, you can try rediscovering the server.

8. Defect SCR 7296 - BFC: Can't create new grid when the checklist shows Replica DB Space Error.

   If the BFC runs out of space before it is able to shutdown, you will need to restart the BFC once you free-up space for it to function properly again.

9. Defect SCR 7312 - Unattended install fails with password !"$%&/()=?'

   If you are performing an unattended install with this version of the product, your password cannot contain a “=”

10. Defect SCR 7376 - STP check is skipped during network detection if server’s public port is configured as trunk

    This bug could occasionally allow servers into grids that should be blocked. If the ports are properly configured, this problem will not be encountered.

11. Defect SCR 7401 - BFC throws "System_limit" error when total no. of characters in the "Edit Grid Parameters" textbox exceeds 256 characters.

    If you need to use more than 256 characters, simply break those parameters into more than one update of the grid.

12. Defect SCR 7413 - BFC UI shows incorrect count of CPU cores when Hyper Threading is disabled.

    On some servers, the CPU count reported by the system is the same when Hyper-Threading is disabled as when it is enabled. This has been observed on some Dell R610s.

13. Defect SCR 7470 - BFC fails to apply grid parameters when more than 1 parameter is passed from the API call.

    This is an issue with how parameters are written to the configuration file passed to aldo set. If in the UI the user enters their data with a comma between entries, the same failure is seen. The work-around for the BFC API is to only pass a single string with a newline separator between the entries.

    For example:

    \"additional_config\":[\"ext_dns1=155.35.34.108\next_dns2=141.202.1.108\"]

    Instead of:

    \"additional_config\":[\"ext_dns1=155.35.34.108\",\"ext_dns2=141.202.1.108\"]

14. Defect SCR 7523 – Upgrade of BFC from 3.1 to 3.5 will fail if a grid contains an application IP address range that was selected from a subnet which was subsequently deleted.

    This issue occurs because deleting a subnet in 3.1 does not properly fail if you have grids with application IP address ranges that are in that subnet. The upgrade process looks for the missing subnet on upgrade and then fails because it is missing. For the workaround, use the instructions from the failed upgrade to restore your previous 3.1 BFC installation. Then, go to each of the grids and remove any application IP address ranges from the grid that does not belong to a currently configured subnet. In some cases, such as when you subsequently re-added the same subnet with a new CIDR prefix length parameter, the range may be within the bounds of a current subnet, but the underlying subnet component will be incorrect and still cause an upgrade failure. You should validate that the subnet in the BFC matches the parameters of the application IP address range in the grid controller UI to be sure.

15. Defect SCR 7047 - Known issue with isotool -o command.

    The isotool -o parameter does not correctly show the USB devices attached to the machine (CentOS 5.5 box). This is a known issue with CentOS 5.5. To resolve this, you must issue the following shell command as user root:

    ```
    service haldaemon restart
    ```

16. Known Issue with Fusion Charts in Internet Explorer 9

    If the graphics rendering option is not set correctly in Internet Explorer 9, graphs in the BFC do not display correctly. Affected graphs appear in the BFC Dashboard, Grids, and Servers pages.

    To fix this problem, in Internet Explorer 9 click Internet Options on the Tools menu. Click the Advanced tab and locate the Accelerated graphics section. Select the Use Software Rendering check box. Save your changes and restart IE.

17. Defect SCR 7707 - During upgrade bfc service hung at starting application 'components' (22 of 45) state.

    This problem can be encountered if you have many VLAN/subnet sets and by-pass the warning during the upgrade that this might be a problem. If you get this warning during the upgrade, please contact Support before proceeding.

18. Defect SCR 7724 - system with 1000 MACs set to AutoDiscovery (blacklist) mode brings BFC down.

    If you need to set these many MACs to AutoDiscovery (blacklist) mode, it would be best to use Manual Configuration (whitelist) mode with 3.5.

19. Defect SCR 7765 - inventory fails when no external IP addresses unavailable.

    Please ensure you have external IP addresses available when added servers to a BFC.

20. Defect SCR 7955 - BMI : Remote Fresh Unattended Installation using PXE Server Failed.

    As part of the process defined for single bare metal ISO creation using the Bare Metal ISO tool, the user needs to place bfcbaremetai.iso and bfcinstall.iso in the directory declared through the 'rawiso' parameter in Bare Metal ISO tool project file. The user then needs to assign '644' permission to these ISOs using following commands:

    chmod 644 bfcbaremetal.iso

    chmod 644 bfcinstall.iso

21. Defect SCR 7984 - Issues with updation of app IP and controller IP after hitting Reset button

    Due to this bug you cannot directly swap the controller and application IPs with each other in one step. If you must do this, first set them to some other values, then you may set them again to the intended values.

22. Defect SCR 8004 - After BFC upgrade from 3.1GA, Administration—Networks—External Tab shows incorrect list of available IPs.

    The IP ranges get mapped to the correct VLANs during an upgrade from pre-3.5.1 to 3.5.1 or later except for one case. If you create a grid and later add a network with a VLAN to that grid, the IPs reserved for that VLAN will remain reserved on the default (untagged) network after the upgrade. If you need to recover those specific IPs, please contact Support.

23. Defect SCR 8005 - API: Issues while adding VLAN to the existing untagged grid via API.

    If you attempt to add a tagged network to an untagged grid via the API, the call will succeed, instead of returning 400 Bad Request.

24. Missing localization strings in BFC 3.5.2.

    Some of the bug fixes and modifications in BFC 3.5.2 changed a small number of strings used in the BFC in 3.5.2. Those modified strings are few in number and mostly in the area of VLAN management, but they will appear in English regardless of your language of choice.

Chapter 4: Key Bug Fixes

The following defects have been resolved in the core CA AppLogic® 3.5 General Availability release:

- 4265: EDT: Safari styling problems
- 4573: EDT: "srv shutdown invalid" causes web shell to abort with error
- 5523: CAT: NAS appliance /etc/resolv.conf file is not being cleanup
- 5561: CAT: All appliance should have a meaningful Instance Name Template

- 5640: VPAT: CA AppLogic® Identifying Cols and rows in data tables criteria issue
- 6127: EDT: Open global grid shell If user not authorized for selected app
- 6629: Product: ESX: grid installation intermittently fails when using an NFS server
- 6678: EDT: Overlapping labels on Chrome 18
- 6713: APK (new-native): pagefile was not deleted after successful installation of Windows APK MSI with automation script.
- 6719: CAT: filer: volume browser help is redirecting to the CA AppLogic® 3.1 document
- 7059: EDT: terminal label text may overlap with terminal images
- 7093: EDT: Modify Boundary window should be adjusted according to screen size (ex: too big to view on laptops)
- 7109: CAT: VPN/VDSes: dns1, dns2 property type is set to string instead of IP
- 7120: WS_API: vdc info method with verbose flag is not working
- 7124: WS_API: Group modify without parameters returns success code instead of failing
- 7125: CLI: Group remove_principal returns success code for non-existing principal instead of failing
- 7130: Filer: Windows: _eXtplorer: volume name displayed as “unknown” in title bar
- 7144: CAT: MYSQLR* web console times out workaround: enable client TCP keepalives
- 7151: Product: MEI: ABS: raw interfaces are not working for windows 2003 appliances
- 7156: APK (new-native): Not able to generate the counter values for CPU summary and volume boot entities works for us, most likely not a bug
- 7159: EDT: No visual indication for drop down list to select VLAN and ip for a raw interface on the Application Configuration window
- 7173: CAT: Windows VDS does not enforce using IP addresses only from the allowed ranges (string vs ip addr prp type)
- 7199: EDT: web shells are closed automatically when number of live web-shells reaches the maximum limit
- 7203: CLI:UTL:wincfg: displays wrong message when component/application is in stopped state
- 7203: Product: wincfg: displays wrong message when component/application is in stopped state
- 7229: EDT: Application start fails silently on unsaved apps

- 7230: EDT: No user prompt when an unsaved application is closed
- 7231: EDT: Appliance status icons may get outside of the appliance box
- 7233: EDT: No button for saving/discarding in ‘Application Start/Restart’ dialog for the application changes before start/Restart of application
- 7234: EDT: PopUp dialog missing when starting the app from UI while vol repair is in progress
- 7238: EDT: Application Start Restart dialog layout is not properly displayed
- 7239: AL grid upgrades failed from 3.0 GA build to 3.5 Betarefresh build
- 7242: EDT: class names are converted to lowercase when dragged to canvas
- 7248: CLI: QTA: user id with description displayed in quota list even after destroyed
- 7255: APK (new-native): UI: Rt.click - Login(ssh) for Windows Filer provisioned application is failing during vol manage with external IP on ESX and Xen grids
- 7258: APK (new-native): IIS base package installation gets stuck during the install and does not work
- 7261: CLI (3tshell.help): The description for --debug and --cap_cpu options needs updating for application start/restart
- 7264: 3tcollect: The log information collected for servers are incorrect when a grid reboots in middle of the execution
- 7266: Product: NWHA: SAN server should be not considered when discovering the network changes "3t grid net_discover" command
- 7269: Product: cannot add a reference server when using a SAN for volume storage
- 7271: CLI (enh): Add support for specifying units when setting vol_repair_speed
- 7272: CLI: Help topics for grid and server set are missing a description of bw_oversubscription parameter
- 7274: Product: removing node from a Xen grid triggers a reboot of the grid controller
- 7275: EDT: Dashboard Network Info, remove "Lan" and change "prefix" to "Netmask"
- 7280: EDT: Pop-up layout is not properly displayed in Chrome19 browser
- 7284: EDT: Round Robin and Round Robin random are working vice versa from UI
- 7285: CLI:MEI:-with 'app config -d' resetting one interface, resets both the interfaces
- 7286: EDT: (enh) Disable Login menu options if application/appliance "not running"
- 7291: Product: sometimes grid info reports Volume Repair Speed as mixed on a fresh build grid

- 7304: EDT: Volume removed from app if app saved after delete volume fails
- 7308: VRM: Controller HA failure on 3 server grid
- 7317: CTL: appliance doesn't restart on the other server when the server was killed
- 7318: MON: IE9: OK button is not working in the New Graph window of MON GUI.
- 7334: CLI:Vol repair speed by default is shown as 9.77 MB/sec
- 7339: VRM: Grid running without controller HA as all 3 servers are in 'none' role
- 7347: EDT: Chrome:PopUp dialog missing when starting the app from UI while vol repair is in progress.
- 7350: EDT:[SAN] Provisioning of a VDS application is failing from UI, whereas from CLI it is working fine.
- 7353: WSAPI: vol create should not require "mirrored" parameter
- 7367: Product: Network interfaces fail on Xen Kernel 3.2.2.11 in version 3.5.12 for Dell 1950/2950 hardware
- 7369: APK (new-native): Automatic Updates are not disabled for Windows 2003 64-bit appliances
- 7370: CAT: VPN appliance client server connectivity is failing
- 7386: CLI: QTA: Wrong output with quota info command in batch mode for available resources
- 7388: Product: NW HA:-The network mode of non-HA/degraded grids is shown as active-passive with HA network disabled/enabled

The following defects have been resolved in the core CA AppLogic® 3.5 beta refresh release:

- 5638: VPAT: CA AppLogic® Text equivalent criteria issue
- 5640: VPAT: CA AppLogic® Identifying Cols and rows in data tables criteria issue.
- 5644: VPAT: CA AppLogic® Language of page criteria issue
- 5955: EDT: FF8: Monitoring graphs do not show on the browser
- 6383: APK (new-native): Issues with Windows automation script that is bundled with APK MSI
- 6511: CAT:-Issues with the BCK Appliance
- 6585: CLI: QTA: Quota limits not enforced except disk immediately after grid reboot
- 6587: CLI: "Put_acl" and "modify_acl" for objects allows multiple owners in the command
- 6652: _eXtplorer : Unable to open eXtplorer(i.e. Volume Browser) for NTFS volumes
- 6663: RBAC & CLI: the can_own attribute of a global group cannot be modified

- 6692: Product: ESX: SAN: controller volume migration from san to local fails
- 6694: CLI: QTA:3t quota modify parent="" does not clear parent
- 6695: CLI: QTA: quotas don't allow user quotas to be modified when maximum level of nested quotas is reached.
- 6696: Product: SAN: The controller restart after local to san migration got failed due to the issue that the loop device that is attached to the metadata volume is being used by another script
- 6704: EDT: Singleton appliance right click view boundary grayed out in view mode
- 6710: CLI: QTA: Dashboard messages for global users and groups not getting displayed
- 6714: WS_API: Invalid XML tags in quota info output
- 6720: EDT: Hidden property value visible when unhidden for an appliance.
- 6727: WS_API: No provision to change quota resource limit to unlimited
- 6735: APK(new-native): fscheck operation is failing with Windows Filer created using Filer_Windows-3.5.4.exe with error 'fsck' mode for filer that is not supported
- 6736: APK(new-native): Files are missing in C:\appliance folder after successful installation of SQL_Windows-3.5.4.exe
- 6737: Product: Running app migration failed when it contains NAS and NASR
- 6739: CLI: 3tctl: volumes migrated from san to local are not mirrored
- 6740: Product: Doc links from appliance GUIs not opening latest docs
- 6742: CLI: QTA: Dashboard messages are not deleted when owners primary group is changed
- 6747: WS_API: Quota list returns error when principal parameter is specified
- 6767: APK(new-native): Shutdown.exe is not copied to SYSWOW64 folder on the WIN0364S and WIN0364E appliances
- 6771: APK(new-native): Color Management settings were in-correct for WIN 2008 (64-bit) Web Edition and the screen(desktop) is not showing up properly
- 6773: APK(new-native): Configure IE Enhanced Security are not disabled for all the editions of 64-bit of WIN 2008 appliances
- 6774: APK(new-native): Windows Updates are not disabled for all the editions of 64-bit of WIN 2008 appliances
- 6790: Product: 3tcollect and 3taudit fail to execute on an upgraded 3.5 grid
- 6808: EDT:-Controller UI logs out when user tries to view/edit the applications (BackupHelper and MigHelper)
- 6837: CLI: QTA: Used resources are not updated after app clean operation

- 6868: EDT: Grid user with valid login access can escalate his permissions
- 6900: CLI: QTA: quota list does not produce output in batch mode when resources have 0 limit
- 6918: EDT: Alert the user that the application changes will take place only when the application is restarted
- 6931: WISA_r11 template not useable, fails start and basic config operations
- 6937: VPAT: CA AppLogic® login Language of page criteria issue
- 6971: CLI: QTA: Exceeded is yes for a User with 0 resource limits
- 6981: CLI: vol creation aborted due to segmentation fault (rare issue)
- 6986: CLI (3tmigrate) - catalog migration fails when executed by a regular user and --preserve_xxx is specified
- 6997: 3tcollect: unable to successfully run 3tcollect on a running grid from the BFC
- 7000: VPAT: screen reader is unable to interpret the Resources graph using the colors
- 7011: EDT: Failed to update package desciptor file when we add and remove MEI interfaces
- 7032: Product: multiple apps can be incorrectly started using the same IP addresses
- 7046: RBAC: Removing a user from a his primary group leaves the group as the user's primary group
- 7052: Product: application migration: "Migrate from" is failing when used in GUI
- 7071: EDT:application migration ("Migrate To") is failing from 3.5.9 to 3.1
- 7088: Product: ESX: "grid set" command with switch ID throws unwanted error message after successful setting the required switch as "active"
- 7123: WS_API:Group put with principal parameter is not working
- 7134: Product: network HA: previously downed NICs are not re-discovered after rebooting the server and restoring the network
- 7137: network ha on external network does not work with 10/100mbps switches.
- 7142: Product: Xen: CA AppLogic® 3.5.7 kernel (v3.2.2) fails to boot on specific hardware when ioatdma is enabled
- 7146: EDT: No alert is displaying from editor window for a locked appliance when we tried to view class descriptor
- 7158: CLI: app config reports syntax error after clearing the iface connection

- 7170: CLI: QTA no dashboard messages displayed when quota for a user is added
- 7171: CLI: QTA: 3t user/group destroy does not delete quota dashboard messages
- 7180: APK (new-native): Windows Filer: boot volume resize operation sometimes fails
- 7201: EDT: Various issues when using FireFox 13
- 7209: APK (new-native): Unable to create NTFS volumes on ESX and Xen grids using the latest version (3.5.10) of APK/EXEs
- 7213: EDT: Chrome unload events not handled properly
- 7235: Product: Controller VM does not migrate when primary server role is changed using "set srv role" command in a 2 node grid
- 7250: EDT: start/restart dialogs should set start button to be the default
- 7251: Product: Security Vulnverability
- 7252: ABS: ESX: changing the network settings incorrectly requires a grid reboot for the changes to take effect

The following defects were resolved in the core CA AppLogic® 3.5 beta initial release:

- SCR 6160: Product: Security vulnerability
- SCR 6235: Product: Notes section in help of hvm2pv utility displays mismatch about state of aplication while executing hvm2pv
- SCR 6406: CPL: Unable to provision same application in different grids at the same time
- SCR 6078: EDT: Chrome: Changes are not saved in the infrastructure editor window
- SCR 6127: EDT: Open global grid shell if user not authorized for selected app
- SCR 6224: EDT: wrong message when try to drag appliance from catalog in view mode
- SCR 6572: EDT: The dialog button not big enough in error prompt dialog
- SCR 6203: CLI: ‘modify group –principal’ does not work when directly using user/group name
- SCR 6204: CLI: Inconsistency between "user" vs. "usr", when name contains non-AlphaNumeric characters
- SCR 6220: CLI: 3tcat import does not clean properly on fail/abort
- SCR 6399: CLI: 3tsrv: does not list the used block devices on an ESX server
- SCR 6563: CLI: perl warnings when importing an app with local catalog
- SCR 6567: CLI: app stop --all fails to stop all the apps if one of the apps is stopped explicitly while the operation is in progress

- SCR 6587: CLI: "Put_acl" and "modify_acl" for objects allows multiple owners in the command
- SCR 6640: CLI: An objects ”put_acl” fails to associate a user as a owner which is recreated
- SCR 6663: CLI: the can_own attribute of a global group cannot be modified
- SCR 5638: VPAT: missing alt text tags for html elements
- SCR 5641: VPAT: Assigned titles to the iframes in Application, Dashboard, logs and Support tabs
- SCR 5644: VPAT: Sets lang attribute on all except doc pages
- SCR 6232: VPAT: Fixed duplicate ID attribute discovered by QA
- SCR 6231: VPAT: Unable to Sort Application List and Navigate Toolbar Buttons Using Keyboard Only
- SCR 6232: VPAT: Parsing criteria - missing open/end tags
- SCR 6233: VPAT: Missing lang attribute
- SCR 6234: VPAT: Screen reader unable to read toolbar buttons and log msgs
- SCR 6140: WS_API: Doesn't work with couple app config parameters: --d mem.min and --template=App1
- SCR 6269: WS_API: WS_API App does not return any value or message for unsupported calls
- SCR 6610: WS_API: Incorrect error code returned for 'group put_newobj_acl' and 'grid put_acl'

The following BFC defects have been resolved in the CA AppLogic® 3.5 release:

- SCR 6069: LDAP loses cached data on BFC restart
- SCR 6131: BFC upgrade stuck at console if yum is locked
- SCR 6219: Race condition in Xen Deployment code which can cause GridOS deployment failures
- SCR 6229: "=" operator does not work for "Discovery Date" property in Edit smart tag window
- SCR 6230: In Edit smart tag window "Value" field for "Total Disk Capacity" property will not work for decimal values
- SCR 6262: Edit smart tag window does not display servers with required tag
- SCR 6275: BFC: Servers are not getting discovered after reinstalling the BFC
- SCR 6298: BFC API: Some wrong information present in output of Grid List Call
- SCR 6337: Incorrect network inventory data in network HA environment

- SCR 6344: Unable to add a new MAC address (es) to the existing MAC list through a text file.
- SCR 6355: Inventory info is not shown for server discovered in quarantined mode
- SCR 6379: Version delete can hang if done immediately after version import
- SCR 6425: BFC upgrade from 3.1.0 to 3.5.0 fails if there are grids that were never started
- SCR 6439: External storage test fails in Admin tab if other mandatory fields are not specified yet
- SCR 6442: Bare Metal: want to have support for USB disk in addition to memory stick
- SCR 6461: Grid servers sometimes failed to boot into Xen grid correctly due to timing issue
- SCR 6490: API: Remove of GNS criteria from a grid is not working
- SCR 6500: "Select All" checkbox not functioning for External and Power network tabs
- SCR 6538: Servers whose backbone connection are not GigE are allowed to go into a grid
- SCR 6554: BMI: Password / Password containing all alphabets, special characters, and numbers are not working
- SCr 6564: API: API call accepts the grid ID update but the grid subsequently fails
- SCR 6608: CentOS boots to grub prompt post BMI if second disk was chosen in disk layout screen
- SCR 6625: Grid creation fails if the license key is mentioned during the grid creation.
- SCR 6641: In the Grid Create wizard on Chrome the buttons on the bottom are cut off
- SCR 6665: Grids properties page-> Xen/VMware tab does not show scrollbar if more tag criteria values are specified.
- SCR 6668: Powering off a server via ssh shutdown can lead to ssh hangs in some instances
- SCR 6688: DR hung at "waiting for client interface..."
- SCR 6722: BFC service fails to start if NFS server hosting the data base replica goes offline and subsequently stop_replication script is run
- SCR 6796: Nnfs mount fails due to leading "/" on path
- SCR 6872: BMI recovery Failed with both NFS Share and USB Drive

- SCR 7005: IP addresses that are pingable by the BFC prevents Grid creation on tagged VLANs
- SCR 7018: Grid controller IPs are not released from untagged VLAN in-use IP list even after grid is deleted
- SCR 7087: BFC service restart takes up to two hours to finish
- SCR 7126: BFC upgrade fails when 3.1 grid is configured with multiple external networks and VLANs
- SCR 7117: Controller VM does not failover to the Secondary server in a 2 Node Grid when the primary Server is killed
- SCR 7040: "Incorrect VMware Servers count after SLA is increased, if no matching servers are found"

The following BFC defects have been resolved in the BFC 3.5.1 release:

- SCR 7625: Backbone with non-24 CIDR causes an issue for PXE
- SCR 7523: If a subnet is deleted in 3.1 which is used for application ips in a grid, upgrade will fail
- SCR 7503: Several values on a grid can't be set via the BFC API (Striping, DNS, Time server)
- SCR 7526: Inventory process incorrectly reports STP status in some cases
- SCR 7527: Scalability: Grid delete with many networks/vlans defined takes a long time
- SCR 7659: Grid monitoring can stack up "grid info" calls into the grid controller
- SCR 7675: VLAN-Server goes to quarantine state when external BFC's switch port and server's switch ports are in "access" and "trunk" modes respectively
- SCR 7312: Unattended install fails with password !"$%&/()=?'

The following BFC defects have been resolved in the BFC 3.5.2 release:

- SCR 7363: VLAN ID 0 is not permitted in BFC 3.5
- SCR 7411: BMI: Invalid check for ownership permissions on NFS Replica DB path during BMI
- SCR 7595: grid start order does not get grid controller started first
- SCR 7614: Server state shows incorrectly as Failed Running after reboot of 44 node grid
- SCR 7629: Server state for 60 node grid shows wrongly as "Joining Grid" due to collector timeout while running "3t grid info"
- SCR 7630: BFC UI unresponsive upon deletion of 10 servers (environment specific)

- SCR 7637: BFC operations are stuck after 30 sequential grid creation jobs are submitted
- SCR 7697: Improve UI performance when displaying external networks & for displaying them in grid create
- SCR 7698: When displaying VLANs, the UI takes too long
- SCR 7702: Grid handling of default VLAN during a BFC upgrade can cause the upgrade to fail
- SCR 7713: BFC start is very slow with large numbers of networks/VLANs
- SCR 7733: VLAN creation slowed by lookup of {subnet,user}_reserved pools
- SCR 7746: Slow performance when creating large numbers of subnets and vlans
- SCR 7747: Modifying grid application IPs needs to scale and perform
- SCR 7753: Grids go to uctl_breached after BFC upgrade from 3.5.0 to 3.5.2
- SCR 7754: Lots of redundant LDAP transactions being created which take a long time to replay during LDAP restore.
- SCR 7760: BFC quarantines servers after intermittent IPMI errors
- SCR 7763: Subnet and VLAN deletes take too long at scale
- SCR 7768: Deletion of 300 VLANs from a running grid gets stuck
- SCR 7770: db_restore does not restore bfcadmin key
- SCR 7776: Deletion of multiple grids caused BFC service restart
- SCR 7778: erlang container gets "too many processes" error at large scale
- SCR 7790: BFC start hangs on networking app
- SCR 7823: BFC: Updating application IPs and controller IP of a grid is not working
- SCR 7825: BFC: Adding 2000 VLANs to multiple grids simultaneously caused BFC service restart
- SCR 7832: Inventory Memory Calculation for older hardware (Super Micro) is not correct
- SCR 7839: BFC: BFC UI is hanged for long time during deletion of grid with 2000 VLANs
- SCR 7843: At scale, grid update app ips operation blocks grids for a long time at BFC start
- SCR 7854: IP ranges grid is slow for select all and other things
- SCR 7885: Upgrading a BFC from 3.0 or 3.1 to 3.5.X can create duplicate grid ip ranges in the property

Chapter 5: Hotfixes

This section describes all of the available hotfixes for this release. Make sure that your grid is updated with the mandatory hotfixes to ensure a properly working grid. If a hotfix requires a grid reboot it will be marked as such.

Mandatory Hotfixes

- hf9467 (see page 92): Resolves both an issue where 'comp start' may result in a controller crash and OpenSSL security vulnerabilities.
- hf9466 : Resolves the metering server collecting and reporting data to the configured metering server. This hotfix modifies files on the grid controller and removes all metering log files currently residing on the grid controller.
- hf9184: (see page 97) Resolves a number of Xen security vulnerabilities.
- hf8368: (see page 99) Resolves an issue where regular users with a valid login can access a CA AppLogic® grid and elevate their permissions.
- hf8975 (see page 100): Resolves several device support, security and stability issues (for new grids).
- hf7596 (see page 105): Resolves an issue in CA AppLogic® 3.5.19 where servers with very large data volumes can fail on reboot or upgrade due to a filesystem check running the server out of memory (SCR 7596).
- hf8096 (see page 105): Resolves the following: SCR 8096: Product: Xen: HVM guest freezes during heavy I/O. Also resolves Xen Security Advisories.
- hf8114 (see page 108): Resolves the following: SCR 8114 - Product: lost access to the grid controller while managing a windows volume and copying a file to it from the impex volume. SCR 7673 - Xen netback driver crashes when receiving large packets.
- hf7791 (see page 111): Upgrades the HP Smart Array Controller driver hpsa to the latest version 3.1.0-7.
- hf7694 (see page 112): Resolves several device support, security and stability issues (for existing grids).

Recommended Hotfixes

- hf8751: (see page 93) Resolves performance issues with relatively new CPUs that support Extended Page Tables (EPT).

Optional Hotfixes

- hf9414 (see page 93): Fixes a guest VM network shutdown that can also be accompanied by a dom0 kernel crash that occurs with *some* network cards that support TCP segmentation offload.
- hf9259: (see page 97) Resolves an issue where grid users cannot access the web shell from the CA AppLogic UI when using Google Chrome web browser (version 28 and above).
- hf9042 (see page 99): Resolves an issue in CA AppLogic® 3.5.19 and 3.7.14 where applications containing a swap volume cannot be migrated (SCR 9042).
- hf8315 (see page 108): Resolves SCR 8315 - Product: SAN ESX grid upgrade failure to CA AppLogic® 3.7.
- hf7781 (see page 108): Resolves SCR 7781 - heartbeat logs fill filesystem on dom0 (which may cause a server to fail to boot).
- hf7906 (see page 119): Resolves an issue where applications open slowly (~2 minutes) in the infrastructure editor when opened by a user who is a member of a large number of groups (SCR 7906).
- e7748 (see page 120): Provides an enhancement to CA AppLogic® where it automatically restarts an appliance when CA AppLogic® loses connection to the VM agent running within the appliance.
- hf7720 (see page 120): Resolves an issue where importing a class/catalog from a URL does not work (SCR 7720).
- e7330 (see page 120): Resolves an issue where migrating applications is slow between grids that are geographically separated (SCR 7330).

Language Pack Hotfixes (Internationalization for the CA AppLogic® grid)

- lp4631: Japanese
- lp4632: French
- lp4633: Chinese
- lp4634: Spanish
- lp4635: German
- lp5501: Italian

Obsolete Hotfixes

These hotfixes are now obsolete with the release of df7588 and hf7694 and should not be used:

- df7588 (see page 115): Resolves several device support, security and stability issues (for new grid installations).
- hf8017 (see page 109): Resolves Xen Security Advisories (XSAs) 6 - 31.
- hf8003 (see page 119): Resolves SCR 7673 - Xen netback driver crashes when receiving large packets.
- hf7759 (see page 112): Resolves Xen Security Advisory XSA-25 http://seclists.org/oss-sec/2012/q4/141, where an Out-of-memory error can occur due to a malicious kernel/ramdisk.
- hf7611 (see page 121): Resolves Xen Security Advisories.
- hf6169 (see page 122): Fixes an issue where Windows appliances fail to start after their hostname has been changed.
- hf7516 (see page 122): Fixes an issue where the link state for an Emulex 10G NIC is wrongly reported as unknown or down due to the bug in the Emulex NIC driver be2net.

Notes:

- You can now control the Native Command Queuing (NCQ) for hard disks by using the grid set command (grid set disk_ncq_enable). This is used for hard disks where NCQ must be turned off due to possible stability issues. In a future release, you will also be able to also control other settings such as bandwidth oversubscription.
- You can now control the volume repair speed (see page 15) using the grid set command (grid set vol_repair_speed).

Hotfix hf9467

Release date: 2014-7-25

Applicability:

hf9467 is a mandatory hotfix for CA AppLogic 3.5.19, 3.7.14 and 3.8.6. This hotfix cannot be used with other versions of CA AppLogic.

The hotfix resolves the following:

- Issue where ‘comp start’ may result in a controller crash
- OpenSSL security vulnerabilities

The hotfix replaces the control domain kernel on grid servers. A grid reboot is not required for the update to take effect.

Hotfix hf9466

Release date: 2014-8-29

Applicability:

hf9466 is a mandatory hotfix for CA AppLogic 3.x and later. This hotfix cannot be used with other versions of CA AppLogic.

The hotfix disables grid metering by modifying files on the grid controller. This causes the controller to stop collecting and reporting metering data to the configured metering server. In addition, this hotfix removes all metering log files currently residing on the grid controller.

Note: This hot fix must be installed on future grids that you will create.

A grid reboot is not required for the update to take effect.

Hotfix hf9414

Release date: 2014-2-19 (3.5 and 3.7), 2014-3-13 (3.8)

Applicability:

hf9414 is an optional hotfix for AppLogic, applicable to all released versions of AppLogic 3.5, 3.7, and 3.8.

The hotfix fixes a potential guest VM network shutdown that occurs with various network cards supporting TCP segmentation. Additionally, the network shutdown could potentially be accompanied by a dom0 kernel crash.

The hotfix replaces the control domain kernel on grid servers. A grid reboot is required for the update to take effect.

Hotfix hf8751

Release date: 2013-10-14

hf8751 resolves the performance issue reported in bug SCR 8751. This bug was originally reported with the title *Performance issues with megaraid_sas driver in Applogic 3.1+* though it was later determined that it is CPU-related, and not related to the Megaraid device.

The issue only occurs on relatively new CPUs. The exact range of CPU models that are affected has not been determined, but all CPUs that support Extended Page Tables (EPT) appear to be affected.

This hotfix requires a grid reboot.

Applicability:

hf8751 is an optional but highly recommended hotfix for CA AppLogic® users that want to gain improved performance on certain newer CPU architectures. The hotfix is released in two different packages, one for AppLogic 3.5 and another for AppLogic 3.7.

Details:

The hotfix addresses a bug in the Linux kernel, which causes it to waste a lot of CPU time during I/O processing when running in PV mode on XEN. The bug affects *both* the AppLogic control OS (XEN Dom0) and guest VMs running a Linux-based OS (DomU).

Note: To obtain the maximum benefit from applying the fix, any guest VMs (appliances) should also have an updated kernel.

This hotfix replaces the control domain (Dom0) kernel and also provides a pre-built guest kernel that can be used in appliances, if necessary.

If your virtual appliance is using a stock OS kernel, the following operating systems are affected by the bug and must have their kernel replaced in addition to installing the hotfix:

- CentOS-6
- Ubuntu 10.4

If your virtual appliance is using a stock OS kernel, the following operating systems are not affected by the bug:

- CentOS-5

  Note: CentOS-5 comes with Linux 2.6.18, and does *not* exhibit the excessive CPU load on I/O when running in XEN. However, this kernel performs overall slightly worse than newer Linux kernels when used in a virtual machine. If performance is important, using a new kernel can help, but note that if a new kernel is installed in place of the CentOS kernel, the automatic update from the CentOS repositories does not keep the kernel up to date.

- Ubuntu 12.4 (if updated to the latest kernel from the apt repository)

If your virtual appliance is using a custom-built kernel, or using one of the pre-built para-virtualized kernels provided with AppLogic:

- linux-2.6.18 is not affected, no action is needed (though replacing it may improve performance as described in the previous Note).
- linux-2.6.30, 2.6.32.x (longterm stable branch) and newer 2.6 kernels are affected and need to be replaced or patched.
- linux-3.2 versions 3.2.46 and later have the bug fixed. Earlier versions need to be replaced.

Other versions of Linux not already mentioned:

- Check if the kernel has the patch found in the mainline kernel archive that is listed in the GIT change log as:
  xen: correctly check for pending events when restoring irq flags

Preparing for Installation of a New Kernel in Appliances with linux-2.6

Important: When installing the new (linux-3.2) kernel into appliances that were using linux-2.6, it may be necessary to modify the boot command line *and* the system mount table. This is because the new kernel uses virtual disk device names that are different from those used by Linux-2.6. Complete the following steps *before* attempting to install the new kernel and do not reboot the appliance without completing the boot configuration and mount table changes.

1. Run the blkid command to display labels and UUIDs of all disks. Either can be used for specifying mounts, UUID is recommended.

   ```
   blkid
   ```

   The output will be similar to this:

   ```
   /dev/hda1: UUID="5090dfea-e6a1-4aed-b758-509b1fd6f07e" TYPE="ext3"
   ```

2. Ensure your operating system is configured to use plain device names by searching for the following strings in the cat /boot/grub/menu.list and cat /etc/fstab files:

   - /dev/hda1
   - /dev/sda1

3. If the boot device or mounts are specified by device name, change them to use UUID.

   For example, if the following line is included in the menu.lst file:

   ```
   kernel  /boot/vmlinuz-2.6.18-194.32.1.el5xen root=/dev/hda1 ro
   ```

   change it as follows (using your device UUID, not the one in this example):

   ```
   kernel  /boot/vmlinuz-2.6.18-194.32.1.el5xen root=root=UUID=5090dfea-e6a1-4aed-b758-509b1fd6f07e ro
   ```

4. Edit the /etc/fstab file to use UUID. For example, change the following line:

   ```
   /dev/hda1   /   ext3   defaults,noatime      0 0
   ```

   to:

   ```
   UUID=5090dfea-e6a1-4aed-b758-509b1fd6f07e  / ext3  defaults,noatime   0 0
   ```

Installing the Guest Kernel in Appliances

For new appliances being created by installing a new Linux-based operating system, you can use the hvm2pv utility to install the new kernel that is included with this hotfix.

For existing appliances, or if you are not using the hvm2pv utility, complete the following steps:

1. Log in to the appliance using a text ssh shell (for example, from the GUI interface, or using the comp ssh command in the CLI).
2. Download the new kernel and a copy of APK (even if APK is already installed) as follows:

   ```
   cd /tmp
   ctl="${SSH_CLIENT%% *}"
   ```

   For 64-bit appliances:

   ```
   wget http://$ctl:8080/download/domu-3.2.0.53.amd64.tar.gz
   ```

   For 32-bit appliances:

   ```
   wget http://$ctl:8080/download/domu-3.2.0.53.i386.tar.gz
   ```

   In the following command, adjust the APK version to whatever your grid has available in the download directory. Also, use the APK version that corresponds with your operating system (for example, apk-*-linux-rh for CentOS and RHEL, # -linux-ub for Ubuntu and Debian, and so on):

   ```
   wget http://$ctl:8080/download/apk-3.0.11-linux-rh.tar.gz
   ```

3. Unpack the archives as follows (this example uses a 32-bit CentOS appliance and apk-3.0.11):

   ```
   cd / tar -zxf tmp/domu-3.2.0.53.i386.tar.gz tar -zxf tmp/apk-3.0.11-linux-rh.tar.gz
   ```

4. Run the APK installation program as follows:

   ```
   tmp/apk-install
   ```

   This configures the new kernel to boot in XEN PV mode.

Hotfix hf9259

Release date: 2013-09-23

hf9259 resolves SCR 9259 where grid users cannot access the web shell from the CA AppLogic UI when using Google Chrome web browser (version 28 and above).

Applicability:

This hotfix is an optional hotfix for CA AppLogic 3.1.14 and 3.5.19.

This hotfix updates files on the controller; for the hotfix to take effect, refresh the browser cache.

Hotfix hf9184

Release date: 2013-8-14

Applicability:

hf9184 is a mandatory hotfix for CA AppLogic 3.5.19 and 3.7.14 (and may not be used with any other CA AppLogic® version).

hf9184 resolves SCR 9184, SCR 8354, SCR 8017, SCR 7611, and SCR 7759 and the following Xen Security Vulnerabilities:

- XSA-6 HVM e1000, buffer overflow.
- XSA-7 64-bit PV guest privilege escalation vulnerability.
- XSA-8 Guest denial of service on syscall/sysenter exception generation.
- XSA-9 PV guest host Denial of Service.
- XSA-11 HVM guest destroy p2m teardown host DoS vulnerability.
- XSA-12 Hypercall set_debugreg vulnerability.
- XSA-13 Hypercall physdev_get_free_pirq vulnerability.
- XSA-14 XENMEM_populate_physmap DoS vulnerability.
- XSA-16 PHYSDEVOP_map_pirq index vulnerability.
- XSA-17 Qemu VT100 emulation vulnerability.
- XSA-19 Guest administrator can access qemu monitor console.
- XSA-20 Timer overflow DoS vulnerability.

- XSA-21 pirq range check DoS vulnerability.
- XSA-22 Memory mapping failure DoS vulnerability.
- XSA-23 Unhooking empty PAE entries DoS vulnerability.
- XSA-24 Grant table hypercall infinite loop DoS vulnerability.
- XSA-25 Xen domain builder Out-of-memory due to malicious kernel/ramdisk.
- XSA-26 Grant table version switch list corruption vulnerability.
- XSA-27 several HVM operations do not validate the range of their inputs.
- XSA-28 HVMOP_get_mem_access crash / HVMOP_set_mem_access information leak.
- XSA-29 XENMEM_exchange may overwrite hypervisor memory.
- XSA-30 Broken error handling in guest_physmap_mark_populate_on_demand().
- XSA-31 Several memory hypercall operations allow invalid extent order values.
- XSA-33 VT-d interrupt remapping source validation flaw.
- XSA 36 interrupt remap entries shared and old ones not cleared on AMD IOMMUs.
- XSA 38 oxenstored incorrect handling of certain Xenbus ring states.
- XSA 39 Linux netback DoS via malicious guest ring.
- XSA 40 Linux stack corruption in xen_failsafe_callback for 32bit PVOPS guests.
- XSA 41 qemu (e1000 device driver): Buffer overflow when processing large packets .
- XSA 42 Linux kernel hits general protection if %ds is corrupt for 32-bit PVOPS.
- XSA 43 Linux pciback DoS via not rate limited log messages.
- XSA 44 Xen PV DoS vulnerability with SYSENTER.
- XSA 46 Several access permission issues with IRQs for unprivileged guests.
- XSA 47 Potential use of freed memory in event channel operations.
- XSA 50 grant table hypercall acquire/release imbalance.
- XSA 54 Hypervisor crash due to missing exception recovery on XSETBV.
- XSA 55 Multiple vulnerabilities in libelf PV kernel handling.
- XSA 56 Buffer overflow in xencontrol Python bindings affecting xend.

Note: For CA AppLogic® 3.5.19, hf9184 includes and obsoletes hf8017, hf7759, and hf7611.

This hotfix replaces files on the Xen-based servers of the grid. Therefore, a grid reboot is required after applying this hotfix.

If this hotfix is installed on a grid and afterwards a new server is added to the grid, this hotfix is automatically installed on the newly added server, however the newly added server must be rebooted for the hotfix to take effect. If a grid is created and hf9184 is installed during the grid creation process, the grid needs to be rebooted after the grid creation has completed in order for the hotfix to take effect.

Hotfix hf8368

Release date: 2013-07-31

hf8368 is a mandatory hotfix for an issue where regular users with a valid login to a CA AppLogic® grid can elevate their permissions. This vulnerability *cannot* be exploited with an invalid login to the grid.

Applicability:

This affects versions 2.9.9, 3.0.30, 3.1.14, and 3.5.19 and should be applied as soon as possible.

For versions 3.0.30 or later, this hotfix can be downloaded using the Backbone Fabric Controller (BFC).

Note: This hotfix does not require a grid reboot after being applied. It takes effect immediately after installation.

Hotfix hf9042

This hotfix resolves an issue in CA AppLogic® 3.5.19 and 3.7.14 where applications containing a swap volume cannot be migrated.

Applicability:

hf9042 is an optional hotfix for CA AppLogic® 3.5.19 and 3.7.14. This hotfix cannot be used with other versions of CA AppLogic®.

The hotfix modifies files on the grid controller. The hotfix takes effect immediately upon successful installation (no reboot is required).

Dependencies: For CA AppLogic® version 3.5.19, hotfix hf8368 must be installed prior to installing this hotfix; there is no such requirement when installing this hotfix on a 3.7.14 grid.

Hotfix df8975

df8975 is a distribution hotfix for CA AppLogic® 3.5.19. df8975 resolves the following issues. Note that all bug fixes affect only Xen-based grids unless otherwise noted.

1. SCR 8975 where dom0 kernel crashes on boot due to buggy CPU tables (Dell C6200 hardware).
2. SCR 7507 where Broadcom 10G NICs are not supported due to missing firmware files for the Broadcom bnx2x driver.
3. SCR 7533 where some of the firmware files are missing for ~50 devices (NICs and hard disk controllers).
4. SCR 7669 where grid creation fails on servers using Broadcom BCM5709 NICs
5. SCR 6779 where Broadcom Gigabit Ethernet NICs sometimes do not operate at Gigabit speed.
6. SCR 7516 where NIC state is wrongly reported for Emulex 10G-based NICs
7. SCR 7528 where Intel e1000e NIC driver hardware hangs
8. SCR 7640 where the Linux kernel crashes on boot when using the HP DL380 G6 disk controller
9. SCR 6448 where Dell PERC H200 controller disks are reported as zero capacity due to a bug in the mpt2sas driver.

In addition, df8975 resolves the following grid stability issues:

1. SCR 8096 where HVM guest freezes during heavy I/O.
2. SCR 7791 which upgrades hpsa driver with respect to advisory: HP Smart Array Controllers – driver upgrade required to prevent a read only Linux File System or an unresponsive server when a target reset is issued by the Smart Array Controller Driver for Linux (hpsa).
3. SCR 8114 where access to the grid controller is lost while managing a windows volume and copying a file to it from the impex volume.
4. SCR 7673 where the Xen netback driver crashes while processing large network packets
5. SCR 7611 resolves several Xen security advisories.
6. SCR 7665 where stopping multiple appliances can overload the server and lead to cleanup failures causing grid instability.
7. SCR 7569 where app stop hangs after applications has been running for more than 10 days (ESX-based grids only).
8. SCR 7666 where an unnecessary error message is printed when a packet is dropped by a network queue.
9. SCR 6169 where Windows appliances fail to start after the hostname is changed.

10. SCR 7658 where not all of the physical disks of the servers are utilized upon grid creation.
11. SCR 7667 where upon a server reboot Xen leaves behind large files that are never deleted.

The resolution to the above SCRs affects grid servers that use any of the following hardware:

1. Broadcom NetXtreme II 1 Gigabit Ethernet controller chipsets BCM5706/BCM5708/5709/5716
2. Broadcom NetXtreme II 10 Gigabit Ethernet controller chipsets BCM57710/BCM57711/BCM57711E/BCM57712
3. ServerEngines Corp. Emulex OneConnect 10Gb NIC (be3) (rev 01)
4. Devices supported by Intel e1000e NIC driver:

- Intel® 82563 Gigabit Ethernet PHY
- Intel® 82566 Gigabit Ethernet PHY
- Intel® 82567 Gigabit Ethernet Controller
- Intel® 82571EB Gigabit Ethernet Controller
- Intel® 82572EI Gigabit Ethernet Controller
- Intel® 82573E Gigabit Ethernet Controller
- Intel® 82573L Gigabit Ethernet Controller
- Intel® 82573V Gigabit Ethernet Controller
- Intel® 82574 Gigabit Ethernet Controller
- Intel® 82577 Gigabit Ethernet PHY
- Intel® 82578 Gigabit Ethernet PHY
- Intel® 82579 Gigabit Ethernet Controller
- Intel® 82583V Gigabit Ethernet Controller
- Intel® Gigabit CT Desktop Adapter
- Intel® Gigabit PT Quad Port Server ExpressModule
- Intel® PRO/1000 PF Dual Port Server Adapter
- Intel® PRO/1000 PF Quad Port Server Adapter
- Intel® PRO/1000 PF Server Adapter
- Intel® PRO/1000 PT Desktop Adapter
- Intel® PRO/1000 PT Dual Port Server Adapter
- Intel® PRO/1000 PT Quad Port Low Profile Server Adapter
- Intel® PRO/1000 PT Quad Port Server Adapter

- Intel® PRO/1000 PT Server Adapter

1. Dell PowerEdge RAID Controller (PERC) H200 disk controller
2. Dell PowerEdge C6200 Server
3. HP DL380 with Smart Array G6 P410i disk controller

This hotfix resolves the following Xen Security Advisories:

XSA-6 http://lists.xen.org/archives/html/xen-devel/2012-02/msg00212.html HVM e1000, buffer overflow.

XSA-7 http://lists.xen.org/archives/html/xen-announce/2012-06/msg00001.html 64-bit PV guest privilege escalation vulnerability.

XSA-8 http://lists.xen.org/archives/html/xen-announce/2012-06/msg00003.html Guest denial of service on syscall/sysenter exception generation.

XSA-9 http://lists.xen.org/archives/html/xen-announce/2012-06/msg00002.html PV guest host Denial of Service.

XSA-11 http://lists.xen.org/archives/html/xen-devel/2012-08/msg00855.html HVM guest destroy p2m teardown host DoS vulnerability.

XSA-12 http://lists.xen.org/archives/html/xen-announce/2012-09/msg00000.html Hypercall set_debugreg vulnerability.

XSA-13 http://lists.xen.org/archives/html/xen-announce/2012-09/msg00001.html Hypercall physdev_get_free_pirq vulnerability.

XSA-14 http://lists.xen.org/archives/html/xen-announce/2012-09/msg00002.html XENMEM_populate_physmap DoS vulnerability.

XSA-16 http://lists.xen.org/archives/html/xen-announce/2012-09/msg00005.html PHYSDEVOP_map_pirq index vulnerability.

XSA-17 http://lists.xen.org/archives/html/xen-announce/2012-09/msg00003.html Qemu VT100 emulation vulnerability.

XSA-19 http://lists.xen.org/archives/html/xen-announce/2012-09/msg00007.html Guest administrator can access qemu monitor console.

XSA-20 http://lists.xen.org/archives/html/xen-announce/2012-11/msg00001.html Timer overflow DoS vulnerability.

XSA-21 http://lists.xen.org/archives/html/xen-announce/2012-11/msg00003.html pirq range check DoS vulnerability.

XSA-22 http://lists.xen.org/archives/html/xen-announce/2012-11/msg00005.html Memory mapping failure DoS vulnerability.

XSA-23 http://lists.xen.org/archives/html/xen-announce/2012-11/msg00004.html Unhooking empty PAE entries DoS vulnerability.

XSA-24 http://lists.xen.org/archives/html/xen-announce/2012-11/msg00002.html Grant table hypercall infinite loop DoS vulnerability.

XSA-25 http://lists.xen.org/archives/html/xen-announce/2012-11/msg00006.html Xen domain builder Out-of-memory due to malicious kernel/ramdisk.

XSA-26 http://lists.xen.org/archives/html/xen-announce/2012-12/msg00000.html Grant table version switch list corruption vulnerability.

XSA-27 http://lists.xen.org/archives/html/xen-announce/2012-12/msg00006.html several HVM operations do not validate the range of their inputs.

XSA-28 http://lists.xen.org/archives/html/xen-announce/2012-12/msg00003.html HVMOP_get_mem_access crash / HVMOP_set_mem_access information leak.

XSA-29 http://lists.xen.org/archives/html/xen-announce/2012-12/msg00004.html XENMEM_exchange may overwrite hypervisor memory.

XSA-30 http://lists.xen.org/archives/html/xen-announce/2012-12/msg00005.html Broken error handling in guest_physmap_mark_populate_on_demand().

XSA-31 http://lists.xen.org/archives/html/xen-announce/2012-12/msg00001.html
Several memory hypercall operations allow invalid extent order values.

Note: df8975 obsoletes distro hotfix df7588. df8975 also includes hf8114, hf8017 and hf7791. If the distro hotfix df8975 is installed there is no need of installing hotfix hf7694.

Note: In contrast to a normal hotfix that is applied to a CA AppLogic® grid after it is created, a distribution hotfix is applied to the CA AppLogic® distribution image used to create grids. Because this is a distribution hotfix, the installation procedure is different than a standard hotfix. Therefore, this hotfix does not need to be re-applied unless the CA AppLogic® 3.5.19 distribution image is re-imported into the BFC.

This hotfix is only for <prodcutname> 3.5.19 distributions.

Installation

After using BFC to download the new hotfix df8975, login as root to the BFC server and do the following:

If you installed the BFC in the default install location (/opt/bfc), run the following command:
/opt/bfc/bin/apply-version-hotfix -v 3.5.19 -f <download directory>/3.5.19/applogic-3.5.19-df8975-part1.2.tar.bz2

If you installed the BFC in an alternate location, run the following command:
<non-standard BFC install location>/apply-version-hotfix -v 3.5.19 -f <download directory>/3.5.19/applogic-3.5.19-df8975-part1.2.tar.bz2

Note: <download directory> is the download directory that was specified during the BFC installation. It is also found in the BFC Administration GUI under the Versions tab.

CA AppLogic Distribution Updates

df8975 makes following updates to the CA AppLogic® distro (BFC imported CA AppLogic® distro directory)

1. df8975 updates the following rpms:
    a. xen/linux-3.2.2-35.xen0.i386.rpm
    b. xen/linux-firmware-3.2.2-35.xen0.i386.rpm
    c. xen/hoop-1.1.12-3.2.2x35.xen0.i386.rpm
    d. xen/sd_filter-3.1.1-3.2.2x35.xen0.i386.rpm
    e. xen/vnp-3.0.6-3.2.2x35.xen0.i386.rpm
    f. xen/xen-4.1.2-11.el5.i386.rpm
    g. VRM rpm in the “xen”, “vmware” and “rpms_common” directories

    For CA AppLogic® 3.5.19 the new vrm is vrm-1.9.37-4.el5.i386.rpm
2. df8975 updates server image files in images/GridOS/<version>/

3. df8975 updates aldo to version 3.5.1.5 - rel-ald-3.5.1.5.tar.bz2

Hotfix hf7596

This hotfix resolves an issue in CA AppLogic® 3.5.19 where servers with very large data volumes can fail on reboot or upgrade due to a filesystem check running the server out of memory (SCR 7596).

hf7596 is a mandatory hotfix for CA AppLogic® 3.5.19. This hotfix cannot be used with other versions of CA AppLogic®. This hotfix modifies files on the grid servers. The hotfix takes effect immediately upon successful installation (no reboot is required).

This hotfix should be re-applied to a 3.5 grid before upgrading to a newer release, in order to ensure that the fsck problem does not happen during the upgrade process.

Hotfix hf8096

Hotfix hf8096 is a mandatory hotfix for CA AppLogic® version 3.5.19. hf8096 resolves following SCR:

- SCR#8096: Product: Xen: HVM guest freezes during heavy I/O.

hf8096 also resolves following Xen security vulnerabilities:

- SCR#8017: Product: Xen Security Vulnerabilities.
- SCR#7759: Product: Xen Security Vulnerabilities (xsa#25).
- SCR#7611: Product: Xen Security Vulnerabilities.

This hotfix resolves the following Xen Security Advisories:

XSA-6 http://lists.xen.org/archives/html/xen-devel/2012-02/msg00212.html HVM e1000, buffer overflow.

XSA-7 http://lists.xen.org/archives/html/xen-announce/2012-06/msg00001.html 64-bit PV guest privilege escalation vulnerability.

XSA-8 http://lists.xen.org/archives/html/xen-announce/2012-06/msg00003.html Guest denial of service on syscall/sysenter exception generation.

XSA-9 http://lists.xen.org/archives/html/xen-announce/2012-06/msg00002.html PV guest host Denial of Service.

XSA-11 http://lists.xen.org/archives/html/xen-devel/2012-08/msg00855.html HVM guest destroy p2m teardown host DoS vulnerability.

XSA-12 http://lists.xen.org/archives/html/xen-announce/2012-09/msg00000.html Hypercall set_debugreg vulnerability.

XSA-13 http://lists.xen.org/archives/html/xen-announce/2012-09/msg00001.html Hypercall physdev_get_free_pirq vulnerability.

XSA-14 http://lists.xen.org/archives/html/xen-announce/2012-09/msg00002.html XENMEM_populate_physmap DoS vulnerability.

XSA-16 http://lists.xen.org/archives/html/xen-announce/2012-09/msg00005.html PHYSDEVOP_map_pirq index vulnerability.

XSA-17 http://lists.xen.org/archives/html/xen-announce/2012-09/msg00003.html Qemu VT100 emulation vulnerability.

XSA-19 http://lists.xen.org/archives/html/xen-announce/2012-09/msg00007.html Guest administrator can access qemu monitor console.

XSA-20 http://lists.xen.org/archives/html/xen-announce/2012-11/msg00001.html Timer overflow DoS vulnerability.

XSA-21 http://lists.xen.org/archives/html/xen-announce/2012-11/msg00003.html pirq range check DoS vulnerability.

XSA-22 http://lists.xen.org/archives/html/xen-announce/2012-11/msg00005.html Memory mapping failure DoS vulnerability.

XSA-23 http://lists.xen.org/archives/html/xen-announce/2012-11/msg00004.html Unhooking empty PAE entries DoS vulnerability.

XSA-24 http://lists.xen.org/archives/html/xen-announce/2012-11/msg00002.html Grant table hypercall infinite loop DoS vulnerability.

XSA-25 http://lists.xen.org/archives/html/xen-announce/2012-11/msg00006.html Xen domain builder Out-of-memory due to malicious kernel/ramdisk.

XSA-26 http://lists.xen.org/archives/html/xen-announce/2012-12/msg00000.html Grant table version switch list corruption vulnerability.

XSA-27 http://lists.xen.org/archives/html/xen-announce/2012-12/msg00006.html several HVM operations do not validate the range of their inputs.

XSA-28 http://lists.xen.org/archives/html/xen-announce/2012-12/msg00003.html HVMOP_get_mem_access crash / HVMOP_set_mem_access information leak.

XSA-29 http://lists.xen.org/archives/html/xen-announce/2012-12/msg00004.html XENMEM_exchange may overwrite hypervisor memory.

XSA-30 http://lists.xen.org/archives/html/xen-announce/2012-12/msg00005.html Broken error handling in guest_physmap_mark_populate_on_demand().

XSA-31 http://lists.xen.org/archives/html/xen-announce/2012-12/msg00001.html Several memory hypercall operations allow invalid extent order values.

Note: hf8096 obsoletes hf8017, hf7759 and hf7611. hf8096 is only for Xen-based grids and does not affect ESX-based grids. This hotfix is only for CA AppLogic® 3.5.19 Xen-based grids (and may not be used with any other CA AppLogic® version). This hotfix replaces files on the servers of the grid. Therefore, a grid reboot is required after applying this hotfix. If this hotfix is installed using BFC, on a grid and afterwards a new server is added to the grid, this hotfix is automatically installed on the newly added server; however the newly added server must be rebooted for the hotfix to take effect. If hf8096 is installed during the grid creation process, the grid needs to be rebooted after the grid creation has completed in order for the hotfix to take effect.

Hotfix hf8315

hf8315 is an optional hotfix for CA AppLogic® 3.5.19. hf8315 resolves the following SCRs:

- SCR 8315 - Product: SAN ESX grid upgrade failure to CA AppLogic® 3.7.x

This hotfix is only for CA AppLogic® 3.5.19 grids (and may not be used with any other version). This bug only affects ESX grids with SAN NFS.

This hotfix patches files on the servers of the grid. Therefore, a grid reboot is not required after applying this hotfix. If this hotfix is installed on a grid and afterwards a new server is added to the grid, this hotfix is automatically installed on the newly added server. The newly added server need not be rebooted for the hotfix to take effect. If a grid is created and hf8315 is installed during the grid creation process, the grid reboot is not required after the grid creation has completed in order for the hotfix to take effect.

Hotfix hf8114

hf8114 is a mandatory hotfix for CA AppLogic® 3.5.19. hf8114 resolves the following SCRs:

- SCR 8114 - Product: lost access to the grid controller while managing a windows volume and copying a file to it from the impex volume
- SCR 7673 - Xen netback driver crashes when receiving large packets

Note: hf8114 includes hf7791 and obsoletes hf8003.

This hotfix is only for CA AppLogic® 3.5.19 grids (and may not be used with any other version). This hotfix replaces files on the servers of the grid. Therefore, a grid reboot is required after applying this hotfix. If this hotfix is installed on a grid and afterwards a new server is added to the grid, this hotfix is automatically installed on the newly added server, however the newly added server must be rebooted for the hotfix to take effect. If a grid is created and hf8114 is installed during the grid creation process, the grid needs to be rebooted after the grid creation has completed in order for the hotfix to take effect.

Hotfix hf7781

hf7781 is an optional hotfix for CA AppLogic® 3.1.14 and 3.5.19. hf7781 resolves SCR 7781: heartbeat logs fill file system on dom0 (which may cause a server to fail to boot)

hf7781 is recommended for grids with 3 or more servers (it is not needed for 1 or 2 server grids).

This hotfix is only for CA AppLogic® 3.1.14 and 3.5.19 grids (and may not be used with any other version). This hotfix updates files on the servers of the grid. Therefore, a grid reboot is required after applying this hotfix. If this hotfix is installed on a grid and afterwards a new server is added to the grid, this hotfix is automatically installed on the newly added server, however the newly added server must be rebooted for the hotfix to take effect. If a grid is created and hf7781 is installed during the grid creation process, the grid needs to be rebooted after the grid creation has completed in order for the hotfix to take effect.

Hotfix hf8017

hf8017 is a mandatory hotfix for CA AppLogic® versions 3.1.14 and 3.5.19.

hf8017 resolves SCR#8017: Product: Xen Security Vulnerabilities (XSA #6 - #31).

hf 8017 also resolves SCR#7759, which resolves Xen Security Advisory XSA-25 http://seclists.org/oss-sec/2012/q4/141, where an Out-of-memory error can occur due to a malicious kernel/ramdisk.

This hotfix resolves the following Xen Security Advisories:

XSA-6 http://lists.xen.org/archives/html/xen-devel/2012-02/msg00212.html HVM e1000, buffer overflow.

XSA-7 http://lists.xen.org/archives/html/xen-announce/2012-06/msg00001.html 64-bit PV guest privilege escalation vulnerability.

XSA-8 http://lists.xen.org/archives/html/xen-announce/2012-06/msg00003.html Guest denial of service on syscall/sysenter exception generation.

XSA-9 http://lists.xen.org/archives/html/xen-announce/2012-06/msg00002.html PV guest host Denial of Service.

XSA-11 http://lists.xen.org/archives/html/xen-devel/2012-08/msg00855.html HVM guest destroy p2m teardown host DoS vulnerability.

XSA-12 http://lists.xen.org/archives/html/xen-announce/2012-09/msg00000.html Hypercall set_debugreg vulnerability.

XSA-13 http://lists.xen.org/archives/html/xen-announce/2012-09/msg00001.html Hypercall physdev_get_free_pirq vulnerability.

XSA-14 http://lists.xen.org/archives/html/xen-announce/2012-09/msg00002.html XENMEM_populate_physmap DoS vulnerability.

XSA-16 http://lists.xen.org/archives/html/xen-announce/2012-09/msg00005.html PHYSDEVOP_map_pirq index vulnerability.

XSA-17 http://lists.xen.org/archives/html/xen-announce/2012-09/msg00003.html Qemu VT100 emulation vulnerability.

XSA-19 http://lists.xen.org/archives/html/xen-announce/2012-09/msg00007.html Guest administrator can access qemu monitor console.

XSA-20 http://lists.xen.org/archives/html/xen-announce/2012-11/msg00001.html Timer overflow DoS vulnerability.

XSA-21 http://lists.xen.org/archives/html/xen-announce/2012-11/msg00003.html pirq range check DoS vulnerability.

XSA-22 http://lists.xen.org/archives/html/xen-announce/2012-11/msg00005.html Memory mapping failure DoS vulnerability.

XSA-23 http://lists.xen.org/archives/html/xen-announce/2012-11/msg00004.html Unhooking empty PAE entries DoS vulnerability.

XSA-24 http://lists.xen.org/archives/html/xen-announce/2012-11/msg00002.html Grant table hypercall infinite loop DoS vulnerability.

XSA-25 http://lists.xen.org/archives/html/xen-announce/2012-11/msg00006.html Xen domain builder Out-of-memory due to malicious kernel/ramdisk.

XSA-26 http://lists.xen.org/archives/html/xen-announce/2012-12/msg00000.html Grant table version switch list corruption vulnerability.

XSA-27 http://lists.xen.org/archives/html/xen-announce/2012-12/msg00006.html several HVM operations do not validate the range of their inputs.

XSA-28 http://lists.xen.org/archives/html/xen-announce/2012-12/msg00003.html HVMOP_get_mem_access crash / HVMOP_set_mem_access information leak.

XSA-29 http://lists.xen.org/archives/html/xen-announce/2012-12/msg00004.html XENMEM_exchange may overwrite hypervisor memory.

XSA-30 http://lists.xen.org/archives/html/xen-announce/2012-12/msg00005.html Broken error handling in guest_physmap_mark_populate_on_demand().

XSA-31 http://lists.xen.org/archives/html/xen-announce/2012-12/msg00001.html
Several memory hypercall operations allow invalid extent order values.

Note: hf8017 includes hf7611 and hf7759 (see page 112). hf8017 is only for Xen-based grids and does not affect ESX-based grids.

This hotfix is only for 3.1.14 and 3.5.19 Xen-based grids (and may not be used with any other version). This hotfix replaces files on the servers of the grid. Therefore, a grid reboot is required after applying this hotfix. For versions 3.1.14 and 3.5.19, if this hotfix is installed using BFC, on a grid and afterwards a new server is added to the grid, this hotfix is automatically installed on the newly added server; however the newly added server must be rebooted for the hotfix to take effect. For versions 3.1.14 and 3.5.19, if a grid is created using BFC and hf8017 is installed during the grid creation process, the grid needs to be rebooted after the grid creation has completed in order for the hotfix to take effect.

Hotfix hf7791

hf7791 upgrades the HP Smart Array Controller driver hpsa to the latest version 3.1.0-7. This is with respect to an advisory: HP Smart Array Controllers – driver upgrade required to prevent a read only Linux File System or an unresponsive server when a target reset is issued by the Smart Array Controller Driver for Linux (hpsa).

hf7791 is a mandatory hotfix only if using the server with the HP Smart Array Controller. hf7791 is for CA AppLogic® 3.1.14 and 3.5.19 when used for Xen-based grids (ESX-based grids are not affected by this hotfix).

This hotfix is for CA AppLogic® 3.1.14 and 3.5.19 Xen-based grids (and may not be used with any other CA AppLogic® version). This hotfix replaces files on the servers of the grid. Therefore, a grid reboot is required after applying this hotfix. If this hotfix is installed on a grid and afterwards a new server is added to the grid, this hotfix is automatically installed on the newly added server, however the newly added server must be rebooted for the hotfix to take effect. If a Xen-based grid is created and hf7791 is installed during the grid creation process, the grid needs to be rebooted after the grid creation has completed in order for the hotfix to take effect.

Hotfix hf7759

hf7759 resolves Xen Security Advisory XSA-25 http://seclists.org/oss-sec/2012/q4/141, where an Out-of-memory error can occur due to a malicious kernel/ramdisk.

hf7759 is a mandatory hotfix for CA AppLogic® 3.5.19 when used for Xen-based grids (ESX-based grids are not affected by this hotfix). hf7759 also includes hf7611 which fixes several Xen Security Advisories. If hf7759 is installed, do not install hf7611.

This hotfix is only for CA AppLogic® 3.5.19 Xen-based grids (and may not be used with any other CA AppLogic® version). This hotfix replaces files on the servers of the grid. Therefore, a grid reboot is required after applying this hotfix. If this hotfix is installed on a grid and afterwards a new server is added to the grid, this hotfix is automatically installed on the newly added server, however the newly added server must be rebooted for the hotfix to take effect. If a Xen-based grid is created and hf7759 is installed during the grid creation process, the grid needs to be rebooted after the grid creation has completed in order for the hotfix to take effect.

Hotfix hf7694

hf7694 is a mandatory hotfix for CA AppLogic® 3.5.19 and 3.5.22. hf7694 resolves the following device support issues.

Note: All bug fixes affect only Xen-based grids unless otherwise noted.

1. SCR 7516 where NIC state is wrongly reported for Emulex 10G-based NICs
2. SCR 7528 where Intel e1000e NIC driver hardware hangs

In addition, hf7694 resolves the following grid stability issues:

1. SCR 7611 resolves several Xen security advisories
2. SCR 7665 where stopping multiple appliances can overload the server and lead to cleanup failures causing grid instability
3. SCR 7673 where the Xen netback driver crashes while processing large network packets
4. SCR 7569 where app stop hangs after applications has been running for more than 10 days (ESX-based grids only)

5. SCR 7666 where an unnecessary error message is printed when a packet is dropped by a network queue
6. SCR 6169 where Windows appliances fail to start after the hostname is changed
7. SCR 7658 where not all of the physical disks of the servers are utilized upon grid creation
8. SCR 7667 where upon a server reboot Xen leaves behind large files that are never deleted

The resolution to the above SCRs affects grid servers that use any of the following hardware:

1. Broadcom NetXtreme II 1 Gigabit Ethernet controller chipsets

   BCM5706/BCM5708/5709/5716
2. Broadcom NetXtreme II 10 Gigabit Ethernet controller chipsets

   BCM57710/BCM57711/BCM57711E/BCM57712
3. ServerEngines Corp. Emulex OneConnect 10Gb NIC (be3) (rev 01)
4. Devices supported by Intel e1000e NIC driver

   Intel® 82563 Gigabit Ethernet PHY

   Intel® 82566 Gigabit Ethernet PHY

   Intel® 82567 Gigabit Ethernet Controller

   Intel® 82571EB Gigabit Ethernet Controller

   Intel® 82572EI Gigabit Ethernet Controller

   Intel® 82573E Gigabit Ethernet Controller

   Intel® 82573L Gigabit Ethernet Controller

   Intel® 82573V Gigabit Ethernet Controller

   Intel® 82574 Gigabit Ethernet Controller

   Intel® 82577 Gigabit Ethernet PHY

   Intel® 82578 Gigabit Ethernet PHY

   Intel® 82579 Gigabit Ethernet Controller

   Intel® 82583V Gigabit Ethernet Controller

   Intel® Gigabit CT Desktop Adapter

   Intel® Gigabit PT Quad Port Server ExpressModule

   Intel® PRO/1000 PF Dual Port Server Adapter

   Intel® PRO/1000 PF Quad Port Server Adapter

   Intel® PRO/1000 PF Server Adapter

Intel® PRO/1000 PT Desktop Adapter

Intel® PRO/1000 PT Dual Port Server Adapter

Intel® PRO/1000 PT Quad Port Low Profile Server Adapter

Intel® PRO/1000 PT Quad Port Server Adapter

Intel® PRO/1000 PT Server Adapter

5. Dell PowerEdge RAID Controller (PERC) H200 disk controller
6. HP DL380 with Smart Array G6 P410i disk controller

This hotfix resolves the following Xen Security Advisories:

1. XSA-6 http://lists.xen.org/archives/html/xen-devel/2012-02/msg00212.html HVM e1000, buffer overflow.
2. XSA-7 http://lists.xen.org/archives/html/xen-announce/2012-06/msg00001.html 64-bit PV guest privilege escalation vulnerability.
3. XSA-8 http://lists.xen.org/archives/html/xen-announce/2012-06/msg00003.html Guest denial of service on syscall/sysenter exception generation.
4. XSA-9 http://lists.xen.org/archives/html/xen-announce/2012-06/msg00002.html PV guest host Denial of Service.
5. XSA-11 http://lists.xen.org/archives/html/xen-devel/2012-08/msg00855.html HVM guest destroy p2m teardown host DoS vulnerability.
6. XSA-12 http://lists.xen.org/archives/html/xen-announce/2012-09/msg00000.html Hypercall set_debugreg vulnerability.
7. XSA-13 http://lists.xen.org/archives/html/xen-announce/2012-09/msg00001.html Hypercall physdev_get_free_pirq vulnerability.
8. XSA-14 http://lists.xen.org/archives/html/xen-announce/2012-09/msg00002.html XENMEM_populate_physmap DoS vulnerability.
9. XSA-16 http://lists.xen.org/archives/html/xen-announce/2012-09/msg00005.html PHYSDEVOP_map_pirq index vulnerability.
10. XSA-17 http://lists.xen.org/archives/html/xen-announce/2012-09/msg00003.html Qemu VT100 emulation vulnerability.
11. XSA-19 http://lists.xen.org/archives/html/xen-announce/2012-09/msg00007.html Guest administrator can access qemu monitor console.

Note: hf7694 includes all previous hotfixes that were released for CA AppLogic® 3.5. hf7694 includes hf7611, hf6169 and hf7516. If the distro hotfix df7588 is installed there is no need of installing hotfix hf7694.

This hotfix is only for CA AppLogic® 3.5.19 and 3.5.22 grids (and may not be used with any other CA AppLogic® version). This hotfix replaces files on the servers of the grid. Therefore, a grid reboot is required after applying this hotfix. If this hotfix is installed on a grid and afterwards a new server is added to the grid, this hotfix is automatically installed on the newly added server, however the newly added server must be rebooted for the hotfix to take effect. If a grid is created and hf7694 is installed during the grid creation process, the grid needs to be rebooted after the grid creation has completed in order for the hotfix to take effect.

Hotfix df7588

df7588 is a distribution hotfix for CA AppLogic® 3.5.19 and 3.5.22. df7588 resolves the following device support issues. Note that all bug fixes affect only Xen-based grids unless otherwise noted.

1. SCR 7507 where Broadcom 10G NICs are not supported due to missing firmware files for the Broadcom bnx2x driver.
2. SCR 7533 where some of the firmware files are missing for ~50 devices (NICs and hard disk controllers).
3. SCR 7669 where grid creation fails on servers using Broadcom BCM5709 NICs.
4. SCR 6779 where Broadcom Gigabit Ethernet NICs sometimes do not operate at Gigabit speed.
5. SCR 7516 where NIC state is wrongly reported for Emulex 10G-based NICs.
6. SCR 7528 where Intel e1000e NIC driver hardware hangs.
7. SCR 7640 where the Linux kernel crashes on boot when using the HP DL380 G6 disk controller.
8. SCR 6448 where Dell PERC H200 controller disks are reported as zero capacity due to a bug in the mpt2sas driver.

In addition, df7588 resolves the following grid stability issues:

1. SCR 7611 resolves several Xen security advisories.
2. SCR 7665 where stopping multiple appliances can overload the server and lead to cleanup failures causing grid instability.
3. SCR 7673 where the Xen netback driver crashes while processing large network packets.
4. SCR 7569 where app stop hangs after applications have been running for more than 10 days (ESX-based grids only).
5. SCR 7666 where an unnecessary error message is printed when a packet is dropped by a network queue.

6. SCR 6169 where Windows appliances fail to start after the hostname is changed.
7. SCR 7658 where not all of the physical disks of the servers are utilized upon grid creation.
8. SCR 7667 where upon a server reboot Xen leaves behind large files that are never deleted.

The resolution to the above SCRs affects grid servers that use any of the following hardware:

1. Broadcom NetXtreme II 1 Gigabit Ethernet controller chipsets

   BCM5706/BCM5708/5709/5716
2. Broadcom NetXtreme II 10 Gigabit Ethernet controller chipsets

   BCM57710/BCM57711/BCM57711E/BCM57712
3. ServerEngines Corp. Emulex OneConnect 10Gb NIC (be3) (rev 01)
4. Devices supported by Intel e1000e NIC driver

   Intel® 82563 Gigabit Ethernet PHY

   Intel® 82566 Gigabit Ethernet PHY

   Intel® 82567 Gigabit Ethernet Controller

   Intel® 82571EB Gigabit Ethernet Controller

   Intel® 82572EI Gigabit Ethernet Controller

   Intel® 82573E Gigabit Ethernet Controller

   Intel® 82573L Gigabit Ethernet Controller

   Intel® 82573V Gigabit Ethernet Controller

   Intel® 82574 Gigabit Ethernet Controller

   Intel® 82577 Gigabit Ethernet PHY

   Intel® 82578 Gigabit Ethernet PHY

   Intel® 82579 Gigabit Ethernet Controller

   Intel® 82583V Gigabit Ethernet Controller

   Intel® Gigabit CT Desktop Adapter

   Intel® Gigabit PT Quad Port Server ExpressModule

   Intel® PRO/1000 PF Dual Port Server Adapter

   Intel® PRO/1000 PF Quad Port Server Adapter

Intel® PRO/1000 PF Server Adapter

Intel® PRO/1000 PT Desktop Adapter

Intel® PRO/1000 PT Dual Port Server Adapter

Intel® PRO/1000 PT Quad Port Low Profile Server Adapter

Intel® PRO/1000 PT Quad Port Server Adapter

Intel® PRO/1000 PT Server Adapter

5. Dell PowerEdge RAID Controller (PERC) H200 disk controller
6. HP DL380 with Smart Array G6 P410i disk controller

This hotfix resolves the following Xen Security Advisories:

1. XSA-6 http://lists.xen.org/archives/html/xen-devel/2012-02/msg00212.html HVM e1000, buffer overflow.
2. XSA-7 http://lists.xen.org/archives/html/xen-announce/2012-06/msg00001.html 64-bit PV guest privilege escalation vulnerability.
3. XSA-8 http://lists.xen.org/archives/html/xen-announce/2012-06/msg00003.html Guest denial of service on syscall/sysenter exception generation.
4. XSA-9 http://lists.xen.org/archives/html/xen-announce/2012-06/msg00002.html PV guest host Denial of Service.
5. XSA-11 http://lists.xen.org/archives/html/xen-devel/2012-08/msg00855.html HVM guest destroy p2m teardown host DoS vulnerability.
6. XSA-12 http://lists.xen.org/archives/html/xen-announce/2012-09/msg00000.html Hypercall set_debugreg vulnerability.
7. XSA-13 http://lists.xen.org/archives/html/xen-announce/2012-09/msg00001.html Hypercall physdev_get_free_pirq vulnerability.
8. XSA-14 http://lists.xen.org/archives/html/xen-announce/2012-09/msg00002.html XENMEM_populate_physmap DoS vulnerability.
9. XSA-16 http://lists.xen.org/archives/html/xen-announce/2012-09/msg00005.html PHYSDEVOP_map_pirq index vulnerability.
10. XSA-17 http://lists.xen.org/archives/html/xen-announce/2012-09/msg00003.html Qemu VT100 emulation vulnerability.
11. XSA-19 http://lists.xen.org/archives/html/xen-announce/2012-09/msg00007.html Guest administrator can access qemu monitor console.

Note: df7588 includes all previous hotfixes that were released for CA AppLogic® 3.5. df7588 includes hf7611, hf6169 and hf7516. If the distro hotfix df7588 is installed there is no need of installing hotfix hf7694.

Note: In contrast to a normal hotfix that is applied to a CA AppLogic® grid after it is created, a distribution hotfix is applied to the CA AppLogic® distribution image used to create grids. Because this is a distribution hotfix, the installation procedure is different than a standard hotfix. Therefore, this hotfix does not need to be re-applied unless the CA AppLogic® 3.5.19 or 3.5.22 distribution image is re-imported into the BFC.

Note: Fixes for SCRs 7528, 7640 and 7658 have passed preliminary testing but final testing is still pending.

This hotfix is only for CA AppLogic® 3.5.19 and 3.5.22 distributions.

Installation

After using BFC to download the new hotfix df7588, login as root to the BFC server and do the following:

If you installed the BFC in the default install location (/opt/bfc), run the following command:

/opt/bfc/bin/apply-version-hotfix -v 3.5.19 -f <download directory>/3.5.19/applogic-3.5.19-df7588-part1.2.tar.bz2

If you installed the BFC in an alternate location, run the following command:

<non-standard BFC install location>/apply-version-hotfix -v 3.5.19 -f <download directory>/3.5.19/applogic-3.5.19-df7588-part1.2.tar.bz2

The same commands from above can be used to install the distro hotfix for 3.5.22 as follows:

/opt/bfc/bin/apply-version-hotfix -v 3.5.22 -f <download directory>/3.5.22/applogic-3.5.22-df7588-part1.2.tar.bz2

Note: <download directory> is the download directory that was specified during the BFC installation. It is also found in the BFC Administration GUI under the Versions tab.

Hotfix hf8003

hf8003 resolves SCR 7673 - Xen netback driver crashes when receiving large packets. This issue is typically observed in grids that are using a 10G backbone network. However it is also possible for this issue to occur in grids that use a 1G backbone network if "gro" (generic receive offload) is enabled on the NICs used by the servers of the grid. This hotfix is only for Xen-based grids.

This hotfix is only for xen based 3.5.19 grids (and may not be used with any other CA AppLogic® version). This hotfix replaces files on the servers of the grid. Therefore, a grid reboot is required after applying this hotfix. If this hotfix is installed on a grid and afterwards a new server is added to the grid, this hotfix is automatically installed on the newly added server, however the newly added server must be rebooted for the hotfix to take effect. If a grid is created and hf8003 is installed during the grid creation process, the grid needs to be rebooted after the grid creation has completed in order for the hotfix to take effect.

Hotfix hf7906

This hotfix resolves an issue where applications open slowly (~2 minutes) in the infrastructure editor when opened by a user who is a member of a large number of groups (SCR 7906). After applying this hotfix, the application load time when opened by such a user should take about 10 seconds.

hf7906 is an optional hotfix for CA AppLogic® 3.5.19. This hotfix cannot be used with other versions of CA AppLogic®.This hotfix modifies files on the grid controller. The hotfix takes effect immediately upon successful installation (no reboot is required).

Hotfix e7748

This hotfix provides an enhancement to CA AppLogic® where it automatically restarts an appliance when CA AppLogic® loses connection to the VM agent running within the appliance. This enhancement works only with appliances that have their Field Engineering Code (FE Code) set to 256. This enhancement is used to automatically restart an appliance in cases where the appliance crashes or for some other reason loses network access (which typically causes the VM agent to lose its connection to CA AppLogic®). CA AppLogic® will not automatically restart the appliance if the FE code is not set to 256.

e7748 is an optional hotfix for CA AppLogic® 3.5.19. This hotfix cannot be used with other versions of CA AppLogic®. This hotfix updates binaries on the physical servers. After installing the hotfix, servers that are updated by the hotfix are marked for reboot (and as such, a dashboard message is created to remind users that the servers need to be rebooted). It is necessary to reboot the marked servers in order for the hotfix to take effect. If this hotfix is installed on a grid and afterwards a new server is added to the grid, this hotfix is automatically installed on the newly added server. However the newly added server must be rebooted for the hotfix to take effect.

Hotfix hf7720

This hotfix resolves an issue where importing a class/catalog from a URL does not work (SCR 7720).

hf7720 is an optional hotfix for CA AppLogic® 3.5.x. This hotfix cannot be used with other versions of CA AppLogic®. This hotfix modifies files on the grid controller. The hotfix takes effect immediately upon successful installation (no reboot is required).

Hotfix e7330

This hotfix resolves an issue where migrating applications is slow between grids that are geographically separated (SCR 7330). This hotfix upgrades the openssh package on the grid controller which includes a bug fix that enhances the performance of application migrations. With e7330 applied, application migrations are typically 3-10 times faster (depending on the transfer speeds between the grids and other networking factors).

e7330 is an optional hotfix for CA AppLogic® 3.0.x, 3.1.x and 3.5.x. This hotfix cannot be used with other versions of CA AppLogic®. This hotfix modifies files on the grid controller. The hotfix takes effect immediately upon successful installation (no reboot is required).

Hotfix hf7611

This hotfix resolves the following Xen Security Advisories:

1. XSA-6 http://lists.xen.org/archives/html/xen-devel/2012-02/msg00212.html XSA-6 HVM e1000, buffer overflow.
2. XSA-7 http://lists.xen.org/archives/html/xen-announce/2012-06/msg00001.html 64-bit PV guest privilege escalation vulnerability.
3. XSA-8 http://lists.xen.org/archives/html/xen-announce/2012-06/msg00003.html Guest denial of service on syscall/sysenter exception generation.
4. XSA-9 http://lists.xen.org/archives/html/xen-announce/2012-06/msg00002.html PV guest host Denial of Service.
5. XSA-11 http://lists.xen.org/archives/html/xen-devel/2012-08/msg00855.html HVM guest destroy p2m teardown host DoS vulnerability.
6. XSA-12 http://lists.xen.org/archives/html/xen-announce/2012-09/msg00000.html Hypercall set_debugreg vulnerability.
7. XSA-13 http://lists.xen.org/archives/html/xen-announce/2012-09/msg00001.html Hypercall physdev_get_free_pirq vulnerability.
8. XSA-14 http://lists.xen.org/archives/html/xen-announce/2012-09/msg00002.html XENMEM_populate_physmap DoS vulnerability.
9. XSA-16 http://lists.xen.org/archives/html/xen-announce/2012-09/msg00005.html PHYSDEVOP_map_pirq index vulnerability.
10. XSA-17 http://lists.xen.org/archives/html/xen-announce/2012-09/msg00003.html Qemu VT100 emulation vulnerability.
11. XSA-19 http://lists.xen.org/archives/html/xen-announce/2012-09/msg00007.html Guest administrator can access qemu monitor console.

hf7611 is a mandatory hotfix for CA AppLogic® 3.0.30/3.1.14/3.5.x. This hotfix replaces files on the servers of the grid. Therefore, a grid reboot is required after applying this hotfix. If this hotfix is installed on a grid and afterward a new server is added to the grid, this hotfix is automatically installed on the newly added server. However, the newly added server must be rebooted for the hotfix to take effect. If a Xen-based grid is created and hf7611 is installed during the grid creation process, the grid needs to be rebooted after the grid creation has completed in order for the hotfix to take effect.

Hotfix hf6169

This hotfix fixes an issue where Windows appliances fail to start after their hostname has been changed. On start, the appliance goes into maintenance mode and then fails to start.

hf6169 is an optional hotfix for CA AppLogic® 3.5.19 and should be installed on grids that have Windows appliances. This hotfix cannot be used with other versions of CA AppLogic®. The hotfix updates binaries on the servers. A grid reboot is required after applying the hotfix. If this hotfix is installed on a grid and afterwards a new server is added to the grid, this hotfix is automatically installed on the newly added server. However the newly added server must be rebooted for the hotfix to take effect.

Hotfix hf7516

hf7516 resolves defect #7516, where the link state for an Emulex 10G NIC is wrongly reported as unknown or down due to the bug in the Emulex NIC driver be2net.

hf7516 is an optional hotfix for CA AppLogic® 3.5.19 and 3.5.22 when used for Xen-based grids (ESX-based grids are not affected by this hotfix).

hf7516 only affects grid servers that use following Emulex NIC model:

ServerEngines Corp. Emulex OneConnect 10Gb NIC (be3) (rev 01)

This hotfix is only for CA AppLogic® 3.5.19 Xen-based grids (and may not be used with any other CA AppLogic® version). This hotfix replaces files on the servers of the grid. Therefore, a grid reboot is required after applying this hotfix. If this hotfix is installed on a grid and afterwards a new server is added to the grid, this hotfix is automatically installed on the newly added server, however the newly added server must be rebooted for the hotfix to take effect. If a Xen-based grid is created and hf7516 is installed during the grid creation process, the grid needs to be rebooted after the grid creation has completed in order for the hotfix to take effect.

Appendix A: Third Party Acknowledgments

Click the following link for copyright and licensing agreement information for third-party software used in CA AppLogic®:

Third-Party Acknowledgments