BIND 9.3.2 Release Notes

HP-UX 11i v1, HP-UX 11i v2, and HP-UX 11i v3

HP Part Number: 839997-003
Published: December 2015
Edition: 5

Legal Notices

Contents

1 BIND 9.3.2 Release Notes

This document discusses the most recent product information pertaining to Berkeley Internet Name Domain (BIND) 9.3.2. It also discusses how to install BIND 9.3.2 on the HP-UX 11i v1 and HP-UX 11i v2 operating systems.

Announcement

BIND is a Berkeley implementation of the Domain Name System (DNS). It is a distributed network information lookup service that maps host names to Internet addresses, and Internet addresses to host names. It also facilitates Internet mail routing by providing a list of hosts that accept mail for other hosts.

BIND 9.3.2 is the latest web upgrade version of BIND. It is available for download at http://h20293.www2.hp.com/.

NOTE: BIND 9.3.2 will not be supported on HP-UX 11i v3 after 01 April 2014.

As of the initial release of HP-UX 11i v3, this product is added to new software bundle. As a result, the product can be updated through an HP-UX 11i v3 Operating Environment Update Release (OEUR). It can also be updated through the Application Release (AR) media, through web-releases, and through patches.

What is in this version?

This version of BIND 9.3.2 for the HP-UX 11i v1, HP-UX 11i v2 and HP-UX 11i v3 operating systems does not include any new feature and contains only defect fixes. For information on the defect fixes, see "Defects fixed in this release" (page 15).

BIND 9.3.2 features

BIND 9.3.2 offers the following features:

DNSSEC implementation based on RFC 4033, 4034, and 4035

Starting with BIND 9.3.2, the Domain Name System Security Extensions (DNSSEC) feature implements the standards specified in RFC 4033 (DNS Security Introduction and Requirements),

4034 (Resource Records for the DNS Security Extensions), and 4035 (Protocol Modifications for the DNS Security Extension). The DNSSEC implementation provides the following new features:

- Signed Zone

  A signed zone contains additional security-related resource records (RRs). Table 1 (page 5) describes additional security-related records in BIND 9.3.2.

  Table 1 Security-Related RRs in a Signed Zone

RR Type	Description
DNS Public Key (`DNSKEY`)	Enables normal DNS resolution and stores public keys. The `DNSKEY` record replaces the `KEY` record.
Resource Record Signature (`RRSIG`)	Stores cryptographically generated digital signatures
Next Secure (`NSEC`)	Enables a security-aware resolver to authenticate a negative reply, for non-existence of name or type, using the same mechanism that is used to authenticate other DNS replies. The `NSEC` record replaces the `NXT` record.
Delegation Signer (`DS`)	Simplifies administrative tasks involved in signing delegations across organizational boundaries

- New DNSSEC options in the `options` statement

  BIND 9.3.2 provides new DNSSEC options in the `options` statement. Table 2 (page 5) lists the new options in the options statement located in the `/etc/named.conf` file.

  Table 2 New DNSSEC Options

Option	Description
`dnssec-enable yes_or_no;`	Enables or disables DNSSEC support. If this option is set to **`yes`**, `named` supports the DNSSEC feature. By default, the DNSSEC feature is not enabled.
`dnssec-lookaside domain trust-anchor domain;`	Provides the validator an alternate method to validate `DNSKEY` records at the top of a zone.
`dnssec-must-be-secure domain yes_or_no;`	Specifies hierarchies that are secure (signed and validated). If this option is set to **`yes`**, `named` accepts answers only if they are secure. If this option is set to **`no`**, `named` applies the standard DNSSEC validation.
`disable-algorithms domain { algorithm; [ algorithm; ] };`	Disables the specified DNSSEC algorithms at and below the specified name. Multiple `disable-algorithms` statements are allowed. However, only the most specific is applied.
`sig-validity-interval number;`	Specifies when the automatically generated DNSSEC signatures expire. The default value is 30 days. The maximum is 3660 days (10 years).

  For more information on the new DNSSEC options, see *named.conf*(1)

- New DNSSEC statement in the `options` statement

  BIND 9.3.2 contains trusted-keys, a new DNSSEC statement in the `options` statement located in `/etc/named.conf` file. The `trusted-keys` statement defines DNSSEC security roots. A security root is defined when the public key for a non-authoritative zone cannot be securely obtained through DNS, either because it is the DNS root zone or because its parent zone is unsigned. When a key is configured as a trusted key, it is treated as if it is validated and is secure. The resolver attempts DNSSEC validation on all DNS data in the subdomains of a security root. The `trusted-keys` statement can contain multiple key entries, each consisting of the key's domain name, flags, protocol, algorithm, and the base-64 representation of the key data. For more information on the `trusted-keys` statement, see *named.conf*(1)

Support for the ip6.arpa domain

BIND 9.3.2 uses the `ip6.arpa` domain for IPv6 forward lookups, instead of the `ip6.int` domain. However, BIND 9.3.2 continues to support the `ip6.int` domain for backward compatibility. BIND 9.3.2 also uses the `ip6.arpa` domain for storing IPv6 addresses in the DNS. The existing queries that perform additional section processing to locate IPv4 addresses are redefined to perform additional section processing on both IPv4 and IPv6 addresses.

The `ip6.arpa` domain is a special domain defined to look up a record given an IPv6 address. This domain provides a method to map an IPv6 address to a host name.

An IPv6 address is represented as a name in the `ip6.arpa` domain by a sequence of nibbles separated by dots with the suffix `.ip6.arpa`. The sequence of nibbles is encoded in reverse order wherein the low-order nibble is encoded first, followed by the next low-order nibble and so on. Each nibble is represented by a hexadecimal digit.

For example, consider the following IPv6 address:

```
4321:0:1:2:3:4:567:89ab
```

The following is the reverse lookup domain name in the `ip6.arpa` domain:

```
b.a.9.8.7.6.5.0.4.0.0.0.3.0.0.0.2.0.0.0.1.0.0.0.0.0.0.0.1.2.3.4.ip6.arpa.
```

New method of listing master servers

Starting with BIND 9.3.2, the `masters` statement provides a list of master name servers that can be included in the `masters` clause of the `zone` statement.

The following is the `masters` statement with the new `masters_list` option, which specifies the `acl` name of the list of master name servers:

```
masters name [port ip_port] {(masters_list | ip_addr [port ip_port]
[key key]); [...]};
```

The `masters_list` option specifies one or more IP addresses of master servers, which the slave can contact to update its copy of the zone. The `masters_list` elements can also be names of other master lists. This list can be used in the `masters` clause in the `zone` statement.

The following is a sample `acl` statement that assigns a symbolic name to an address match list:

```
acl acl1 {
         15.70.190.186; 15.70.190.115;
         };
```

The following is a sample `zone` statement with the `masters` clause:

```
zone "example.com" {
type slave;
masters {acl1;};
file "db.example";
};
```

Where:

`acl1` specifies the name of the list of master name servers.

New options in the Options statement

Table 3 lists the new options added in the `options` statement.

Table 3 New Options in the Options Statement

Option	Description
hostname	Identifies the host name of the anycast `named` server that answers the query
`server-id`	Identifies the server ID of the anycast `named` server that answers the query

Table 3 New Options in the Options Statement *(continued)*

Option	Description
key-directory	Specifies the location of the public and private key files if the current directory is not the working directory
memstatistics-file	Specifies the pathname of the file where the server writes memory usage statistics upon exit. The default file is named.memstats.
flush-zones-on-shutdown	Specifies whether any pending zone writes must be flushed when the name server exits because of a SIGTERM signal. The default value is **no**.
check-names	Specifies the list of IPv4 and IPv6 UDP ports that are not used as system assigned source ports for UDP sockets. The default value depends on the usage area. For master zones, the default value is fail. For slave zones, the default value is warn. For an answer (response) received from the network, the default value is ignore.
avoid-v4-udp-ports and avoid-v6-udp-ports	Avoids named from selecting certain ports
use-v4-udp-ports and use-v6-udp-ports	Specifies the port range to be selected by named
query-source-v6	Specifies the address and port used for queries
tcp-listen-queue	This option specifies the length of the listen queue. The default and minimum values are 3. If the kernel supports the dataready accept filter, this option also controls the number of TCP connections that are queued in the kernel space waiting for data, before data is passed to the accept filter.
alt-transfer-source	Specifies an alternate transfer source, if the transfer source listed in the transfer-source option fails and the use-alt-transfer-source option is set.
alt-transfer-source-v6	Specifies an alternate transfer source, if the transfer source listed in the transfer-source-v6 option fails and the use-alt-transfer-source option is set.
use-alt-transfer-source	Specifies whether named must use the alternate transfer sources. Alternate transfer sources are used if views are specified; otherwise, the alternate transfer sources are not used for BIND 8 compatibility.
max-journal-size	Sets a maximum size for each journal file. When the journal file approaches the specified size, older transactions in the journal are removed. The default value is unlimited.
rrset-order	Configures the ordering of records in a multiple record response
preferred-glue	Specifies the glue that is emitted first in the additional section of a query response. If specified, the listed type (A or AAAA) is emitted before any other glue. The default value is NONE if a preference is not set for any type of glue.
root-delegation-only	Switches on the enforcement of delegation-only in top level domains (TLDs) and root zones with an optional exclude list.
querylog	Specifies whether query logging must be started when named starts. If querylog is not specified, query logging is determined by the presence of the logging category queries.
disable-algorithms	Disables the DNSSEC algorithms at and below the specified name. Multiple disable-algorithms statements are allowed. However, only the most specific disable-algorithms option is applied.
max-recursion-depth	To set a limit on the number of levels of recursion named will allow. The default is 7 levels.
max-recursion-queries	To set a limit on the number of iterative queries named will send before terminating a recursive query. The default is 50 queries.

New option to configure the ordering of records

The new `rrset-order` option in the `options` statement enables you to configure the ordering of the records in a multiple-record response. When the name server returns multiple records in a response, it is useful to configure the order of the records placed into the response.

The following is the syntax of the `rrset-order` option:

```
rrset-order {order_spec};
```

Where, an *`order_spec`* can be defined as follows:

```
[class class_name]
[ type type_name ]
[ name domain_name]
order ordering
```

The default value for `class` and `type` is `ANY`, and for `name` is `*`.

The valid values for *`ordering`* are:

`fixed` Records are returned in the order they are defined in the `zone` file

`random` Records are returned in a random order

`cyclic` Records are returned in a round-robin order

The following is an example of the `rrset-order` option:

```
rrset-order {
        class IN type A name "host.example.com" order random;
         order cyclic;
};
```

This `rrset-order` option causes responses for type `A` records in class `IN` that have `host.example.com` as a suffix, to be returned in random order. Other types of records are returned in cyclic order.

If the `options` statement contains multiple `rrset-order` options, they are not combined but only the last `rrset-order` option is used.

New option to set the advertized EDNS UDP buffer size

The `edns-udp-size` option in the `options` statement sets the advertised Extended DNS (EDNS) User Datagram Protocol (UDP) buffer size to enable UDP answers to pass through broken firewalls that block fragmented packets greater than 512 bytes. The valid range of values is 512 to 4096 bytes (values not in this range are adjusted appropriately). The default value of this option is 4096 bytes.

New option to restrict the character set of domain names

This `check-names` option in the `options` statement restricts the character set and syntax of certain domain names in the master files and DNS responses. The rules for valid host names or mail domains are derived from RFC 952 (DoD Internet Host Table Specification) and RFC 821 (Simple Mail Transfer Protocol) as modified by RFC 1123 (Requirements for Internet Hosts - Application and Support). The `check-names` option checks the names of the owner names of `A`, `AAAA`, and `MX` records and also checks domain names in the `RDATA` of `NS`, `SOA`, and `MX` records. It also applies to the `RDATA` of `PTR` records where the owner name indicates that it is a reverse lookup of a hostname (the owner name ends with `in-addr.arpa`, `ip6.arpa`, or `ip6.int`).

The default value of the `check-names` option depends on the usage area. For master zones, the default value is `fail`. For slave zones, the default value is `warn`. For an answer (response) received from the network, the default value is `ignore`.

New options to enable and disable IXFR

In BIND 9.3.2, the incremental zone transfer (IXFR) feature is enabled by default. describes the new options available in the `options` statement that can be used to enable and disable IXFR.

Table 4 Options to Enable and Disable IXFR

Option	Description
`provide-ixfr yes_or_no;`	Determines whether the local server, which acts as a master, responds with an incremental zone transfer when the remove slave server requests an IXFR. If the `provide-ixfr` option is set to **`yes`**, incremental transfer is provided whenever possible. If this option is set to **`no`**, all transfers to the remote server is non-incremental. If the `provide-ixfr` option is not set, the value of `provide-ixfr` in the `view` or `global` options statement is used as default.
`request-ixfr yes_or_no;`	Determines whether the local server, acting as a slave, requests incremental zone transfers from a remote master server. If this option is not set, the value of `request-ixfr` in the `view` or global `options` statement is used as default. If this option is set to **`yes`**, the server, by default, collects statistical data of all zones in the server. If this option is set to **`no`**, the server requests a full zone transfer (AXFR).
`ixfr-from-differences yes_or_no;`	Loads a new version of the master zone from the zone file of the server, or receives a new version of the slave file by a non-incremental zone transfer. If this option is set to **`yes`**, when the server receives a new version of a slave file by a non-incremental zone transfer, the server compares the new version of the master zone with the previous version of master zone and calculates the set of differences. The differences are logged in the journal file of the zone such that the changes can be transmitted to downstream slaves as an incremental zone transfer. If this option is set to **`no`**, the name server must perform a complete zone transfer to the slave server.

Transition support for IPv4 and IPv6

BIND 9.3.2 provides transition support for IPv4 and IPv6 to solve the problem caused by lack of support for either IPv4 or IPv6 address on a host system. It also provides the `dual-stack-servers` option to enable the transition support for IPv4 and IPv6 addresses. This option specifies host names or addresses of systems that access both IPv4 and IPv6 transports. If the host name is specified, a name server must be able to resolve a host name by using only the transport supported by the name server. If the `dual-stack-servers` option is used in dual-stacked system, this option does not have any influence if access to the IPv4 or IPv6 transport is disabled on the command line using the `named -4` command or `named -6` command, respectively.

The syntax for the `dual-stack-servers` option in the `options` statement in the `/etc/named.conf` file is as follows:

```
[ dual-stack-servers [port ip_port] { ( domain_name [port ip_port] |
ip_addr [port ip_port] ) ; ... }; ]
```

New commands in the rndc utility

The following are new commands in the `remote name daemon control (rndc)` utility:

- `retransfer zone [class [view]]`

  This command enables you to retransfer the given zone from the master name server.

- `freeze zone [class [view]]`

  This command enables you to suspend updates to a dynamic zone and enables you to edit a zone that is usually updated dynamically. This command results in changes to the journal

file to be synchronized into the master, and the journal file to be removed. All dynamic update attempts are refused if the zone is frozen.

- `thaw zone [class [view]]`

  This command enables you to update a frozen dynamic zone. This command causes the server to reload the zone from the disk and re-enables dynamic updates after the load is complete.

For more information on these commands, see *rndc*(1). A sample `rndc.conf` file is distributed with this release of BIND in the `/usr/examples/bind` directory. This file can be generated automatically using the `rndc-confgen` utility, which is also distributed with BIND 9.3.2.

New option in the zone statement

The `delegation-only` option is added to the `zone` statement. You can use this option to enforce the delegation-only status of infrastructure zones (for example, `COM`, `NET`, and `ORG`). Any answer that a name server receives without an explicit or implicit delegation in the authority section is treated as `NXDOMAIN`, which indicates that a host name is not found. The `NXDOMAIN` response is the type of response sent by the name server.

New command-line options

Table 5 lists the new command-line options for the various binaries and tools in BIND 9.3.2.

Table 5 New Command-Line Options

Binaries/Tools	Options	Description
`dnssec-keygen`	`-f flag`	Sets the specified flag in the flag field of the `KEY` or `DNSKEY` record. The only recognized flag is Signed Key (`KSK`) `DNSKEY`.
`dnssec-keygen`	`-k`	Generates `KEY` records, instead of the `DNSKEY` records
`dnssec-signzone`	`-g`	Generates `DS` records for child zones from the keyset files. Existing DS records are removed from the signed `db` files.
`dnssec-signzone`	`-k key`	Treats the specified key as a key signing key and ignores any key flags. This option can be specified multiple times.
`dnssec-signzone`	`-l domain`	Generates a DNSSEC lookaside validation (DLV) set in addition to the key (DNSKEY) and `DS` sets. The domain is appended to the name of the records.
`named-checkconf`	`-z`	Performs a check load on the master zone files in the `/etc/named.conf` file
`named-checkconf`	`-j`	Reads the journal while loading a zone file
`named-checkzone`	`-j`	Reads the journal while loading a zone file
`named-checkzone`	`-k mode`	Performs `check-name` checks with the specified failure mode. The values for the failure modes are `fail`, `warn`, and `ignore`. The default value is `warn`.
`named-checkzone`	`-n mode`	Specifies if name server (NS) records must be checked to verify whether they are addresses. The values for this option are `fail`, `warn`, and `ignore`. The default value is `warn`.
`named-checkzone`	`-o filename`	Writes the zone output to the directory
`named-checkzone`	`-t directory`	Specifies the directory under which the `named-checkzone` command is chrooted. The `$INCLUDE` directives in the configuration file are also processed as if they are run by a similarly chrooted `named`.
`named-checkzone`	`-w directory`	Specifies `named` to change to `directory` so that relative filenames in the master file `$INCLUDE` directives are functional. This option is similar to the `directory` clause in the `/etc/named.conf` file.

Table 5 New Command-Line Options *(continued)*

Binaries/Tools	Options	Description
named-checkzone	-D	Specifies the dump zone file in canonical format
named	-4	Specifies named to use only the IPv4 transport even if the host system is capable of handling IPv6 addresses
named	-6	Specifies named to use only the IPv6 transport even if the host system is capable of handling IPv4 addresses
nsupdate	-t	Sets the maximum timeout value for an update request before it can abort. The default value is 300 seconds. To disable the timeout, set this option to 0.
nsupdate	-u	Sets the UDP retry interval. The default value is 3 seconds. If this option is set to 0, the interval is computed from the timeout interval and the number of UDP retries.
nsupdate	-r	Sets the number of UDP retries. The default value is 3. If this option is set to 0, only one update request is made.

Supports RFC 4193 (Unique local IPv6 unicast addresses)

BIND 9.3.2 (C.9.3.2.5.0) for the HP-UX 11i v3 operating system conforms to RFC 4193 (*Unique Local IPv6 Unicast Addresses*). RFC 4193 defines a format for the unique local IPv6 unicast address that is globally unique and not intended for external networks. When named receives an unique local IPv6 unicast address for resolution, it does not send this address to the global DNS server for resolution. Instead, it returns the NXDOMAIN response message by default. As a result, the unique local IPv6 unicast addresses are never exposed to the outside network and are not accessible by external systems.

Changed features

The following are the changed features in BIND 9.3.2:

- In BIND 9.3.2, *named*(1M) selects the best forwarder from the list of forwarders specified in the /etc/named.conf file and sends the query to the forwader with the lowest roundtrip time. In BIND 9.2.0, *named*(1M) does not select a forwarder from the /etc/named.conf file but sequentially sends queries to all the forwarders in the /etc/named.conf file until the query is answered.
- The following DNSSEC features are modified in BIND 9.3.2:
  - In BIND 9.2.0, when the dnssec-keygen command is executed twice with the HMAC-MD5 algorithm, two different key-file pairs are generated. In BIND 9.3.2, the key files are overwritten, resulting in one key-file pair only.
  - In the previous version of BIND, the dnssec-keygen command used the RSAMD5, DH, DSA, RSA, or HMAC-MD5 algorithm. In BIND 9.3.2, the dnssec-keygen command supports only RSASHA1 and DSA algorithms for DNSSEC. HMAC-MD5 and DH are also supported, in which case a KEY record is generated instead of a DNSKEY record. The -k option must be used to generate a KEY record.
  - In BIND 9.3.2, the key file supplied to nsupdate using the -k option must contain a key of the type KEY and not DNSKEY.
  - The dnssec-signzone command creates the db.<*zone*>.signed file, which contains the NSEC (corresponding to the NXT record in 9.2.0) and RRSIG (corresponding to the SIG record in 9.2.0) records. Additionally, it creates a dsset-<*zone*> file that contains the DS record and the keyset-<*zone*> file that contains the DNSKEY record.

- The following `dig` features are modified in BIND 9.3.2:
  - The `-i` option in the `dig` command must be used for `IP6.INT` IPv6 reverse lookups. By default, `dig` performs `IP6.ARPA` reverse IPv6 lookups.
  - The output of the `dig` name command for `Not Implemented` is changed from `NOTIMPL` to `NOTIMP`.
- Table 6 lists the changed command-line options for the `dnssec-signzone` tool in BIND 9.3.2.

Table 6 New Command-Line Options

Binaries/Tools	Old Option	New Option	Changed Functionality
`dnssec-signzone`	`-c cycle-time`	`-c class`	Specifies the DNS class of the zone
`dnssec-signzone`	`-n ncpus`	`-n threads`	No change in the functionality for this option

Installing BIND 9.3.2

This section describes how to install BIND 9.3.2. It also lists the prerequisites for installing BIND 9.3.2.

Prerequisites

Table 7 lists the prerequisites for installing BIND 9.3.2 on the HP-UX 11i v1 and v2 operating systems.

Table 7 BIND 9.3.2 Prerequisites

Operating System	Prerequisite
HP-UX 11i v1	For using DNSSEC Public Key Cryptography functionality, the OpenSSL library must be installed. However, `named` will continue to run without the OpenSSL library.[1]
HP-UX 11i v2	
HP-UX 11i v3	No prerequisites

[1] For the HP-UX 11i v1 operating system, install the OpenSSL software from http://www.software.hp.com to obtain the OpenSSL libraries. For the HP-UX 11i v2 operating system, the OpenSSL libraries are available as part of the core operating system.

NOTE: If you have installed the Web upgrade version of BIND 9.2.0 on an HP-UX 11i v1 system, ensure that you remove the BIND 9.2.0 depot before installing BIND 9.3.2.

Installation instructions

To install BIND 9.3.2, complete the following steps:

1. Review to ensure that your system meets BIND 9.3.2 installation requirements.
2. Go to the HP Software Depot website at:
   http://h20293.www2.hp.com/.
3. Use the **Search** button to browse for BIND. The product catalog page is displayed.
4. Select **BIND** in the product catalog. The BIND page is displayed.
5. Read the "Overview" and "Installation" pages for BIND.
6. Click on **Select** option at the bottom right of any of these pages.
7. Select the appropriate release of HP-UX operating system.
8. Enter the registration information. Read and accept the terms and conditions statements.
9. Click **Next>>**. The Electronic Delivery Receipt page is displayed.

10. Select the BIND 9.3.2 depot under Download Software.
11. Save the BIND 9.3.2 depot in a local directory, for example, `/tmp`.
12. To verify that the BIND 9.3.2 depot is downloaded properly in the local directory, enter the following HP-UX MD5 Secure Checksum command at the HP-UX prompt:

    # **md5sum *<depot_name>***

    The result of this command must match the fingerprint provided in the Electronic Delivery Receipt. If the result does not match, download the BIND 9.3.2 depot again.

    NOTE: The HP-UX MD5 Secure Checksum software is not installed by default on the system. It is available at:

    http://h20293.www2.hp.com/

13. To install the BIND 9.3.2 depot, enter the following command at the HP-UX prompt:# **swinstall -s *<fully_qualified_depot_source_path>***

    The `swinstall` window is displayed.

    NOTE: Ensure that the `DNSUPGRADE.PHNE_33766` product or the `DNSUPGRADE.PHNE_34226` product is installed before installing the `DNSUPGRADE.BindUpgrade` product. If you have installed the PHNE_33766 or PHNE_34226 patch or any of its superseding patches, you need not install the `DNSUPGRADE.PHNE_33766` or `DNSUPGRADE.PHNE_34226` product.

14. Press the space bar to select the product that you wish to install.

IMPORTANT: Do not install Web release versions of BIND prior to BIND 9.3.2, after installing the `DNSUPGRADE.BindUgrade` product. Do not install BIND 9.3.2 version after installing HPUX-NameServer product of BIND 9.7.3 on 11iv3

15. Select **Install** in the Action menu. The Install Analysis window is displayed.
16. Select **OK** when the Status field displays a Ready message. The Install window is displayed.

    The BIND 9.3.2 software installation starts. The `swinstall` command loads the BIND 9.3.2 files on to the system in approximately 3 to 5 minutes.
17. Select **Done** when the Status field displays a Completed message.
18. Select **File->Exit** to exit from the `swinstall` window.

    The `named` daemon is preconfigured and starts after installation. The `swinstall` command installs BIND in the `/opt` directory.

For more information on configuring and using BIND, see the *HP-UX IP Address and Client Management Services Administrator's Guide* at:

http://www.hp.com/go/hpux-networking-docs-11iv3.

Verifying the BIND 9.3.2 installation

To verify whether the BIND 9.3.2 depot is installed successfully on your system, enter the following command at the HP-UX prompt:

swlist -l product *<depot_name>*

If BIND 9.3.2 is installed properly, the following output is displayed:

- On an HP-UX 11i v1 operating system

```
# Initializing...
# Contacting target "hostname"...
#
# Target: hostname:/
#
```

```
BindUpgrade C.9.3.2.N.0 BIND special release upgrade
```

- On an HP-UX 11i v2 operating system

```
# Initializing...
# Contacting target "hostname"...
#
# Target: hostname:/
#

BindUpgrade C.9.3.2.N.0 BIND special release upgrade
```

NOTE:

1. For HP-UX 11i v2 and HP-UX 11i v3, It is admin's role to check that the most recent version of /etc/rc.config.d/namesvrs_dns is in usage, which is located in /usr/newconfig path as /usr/newconfig/etc/rc.config.d/namesvrs_dns.
2. namesvrs_dns script realtime usage should 'always' be from the /etc/rc.config.d path.

Unsupported features

The following are the unsupported features in BIND 9.3.2:

- The following BIND 9.2.0 options are not supported in BIND 9.3.2:
  - The `allow-v6-synthesis` option in `/etc/named.conf` file
  - The `dnssec-makekeyset` command
  - The `dnssec-signkey` command
- The mail destination (`MD`) and mail forwarder (`MF`) records are obsoleted in BIND 9.3.2. The appropriate mail exchanger (`MX`) record must be used in the database (`db`) files.

Known problems

The following are known problems in BIND 9.3.2:

- The DNSSEC public key cryptography in BIND 9.3.2 is not backward compatible. If security is enabled for a zone using the previous DNSSEC feature, the zone must be reconfigured for the new DNSSEC feature by generating keys, signing the zone with the newly generated key, and distributing the new public key.
- The DNSSEC functionality in BIND 9.3.2 is tested only with OpenSSL version A.00.09.07. This functionality may differ with other versions of OpenSSL.

Related information

The following sections discuss the documentation available for BIND 9.3.2.

Manpages

Table 8 describes the manpages distributed with the BIND 9.3.2 depot.

Table 8 BIND 9.3.2 Manpages

Manpage	Description
dnssec-keygen(1)	Tool to generate keys for DNSSEC
dnssec-signzone(1)	Tool to sign the DNSSEC zone
host(1)	Utility for DNS lookup
named-checkconf(1)	Tool to check the syntax of the `named` configuration file

Table 8 BIND 9.3.2 Manpages *(continued)*

Manpage	Description
named-checkzone(1)	Tool to check the validity of a zone
`nslookup.1`*nslookup*(1)	Interactive tool to query name servers
nsupdate(1)	Utility to update the DNS dynamically
rndc-confgen(1)	Tool to generate the `rndc` key
rndc(1)	Utility to control the name server control
dig(1)	Tool to interrogate DNS servers
hosts_to_na(1M)	Command to translate host table to name server file format
lwresd(1M)	Daemon to provide name lookup services to clients that use the BIND 9 lightweight resolver library
named(1M)	Daemon that reads the BIND configuration file, `/etc/named.conf` for initial data on resource records, and listens for queries. The named daemon is the Internet domain name server, and it requires superuser privileges to execute.
sig_named(1M)	Daemon that send signals to the domain name server
named.conf(4)	Configuration file for name daemon
rndc.conf(4)	Configuration file rndc

The `nslookup(1)`, `dig(1M)`, and `host(1)` can be used to troubleshoot BIND 9.3.2. For detailed information and examples of utilities and commands listed in Table 8, see the respective manpages.

Defects fixed in this release

This section discusses the defects fixed in the HP-UX 11i v1, HP-UX 11i v2, and HP-UX 11i v3 operating systems.

It discusses the following topics:

Defects fixed in the HP-UX 11i v1 and HP-UX 11i v2 operating systems

Table 9 lists the defects fixed in BIND 9.3.2 in both the HP-UX 11i v1 and HP-UX 11i v2 operating systems.

Table 9 Defects Fixed in the HP-UX in HP-UX 11i v1 and HP-UX 11i v2 operating systems

Defects fixed in BIND 9.3.2(C.9.3.2.15.0)	
QXCR1001455249 (CVE-2015-8000) (Only on 11iv1)	`named(1M)` accepts some records with an incorrect class instead of rejecting them as malformed due to an error in the parsing of incoming responses.
QXCR1001440768 (CVE-2015-5722)	BIND servers performing validation on DNSSEC-signed records can be exploited remotely causing denial-of-service. Validating recursive resolvers are at greater risk from this defect.
QXCR1001445978 (Only on 11iv2)	`named(1M)` exits due to an assertion failure in `resolver.c` file.
QXCR1001453568 (Only on 11iv1)	`named(1M)` runs out of recursive clients and then hangs.
Defects fixed in BIND 9.3.2(C.9.3.2.14.0)	
QXCR1001389999(CVE-2014-8500)	A flaw in delegation handling could be exploited to put named into an infinite loop, in which each lookup of a name server triggered additional lookups of more name servers.
QXCR1001432300(CVE-2015-5477)	An error in the handling of TKEY queries can be exploited by an attacker for use as a denial-of-service vector, as a constructed packet can use the defect to trigger a REQUIRE assertion failure, causing BIND to exit.
QXCR1001432288	BIND under HP-UX 11.23 can still crash in `pthread_mutex_destroy()`.
Defects fixed in BIND 9.3.2(C.9.3.2.13.0)	
QXCR1001251264	It is a design limitation in named(1M) due to which the information can be prolonged in the cache beyond the period supposedly allowed by the TTL value, causing named(1M) to potentially return incorrect answers.
QXCR1001252460 (Only on 11iv2)	It is a limitation in control scripts of named(1m) due to which named.64 is not getting killed.
QXCR1001252487 (Only on 11iv2)	Control scripts doesn't handle the scenario where a namesvrs_dns script is already existing in /etc/rc.config.d path.
Defects fixed in BIND 9.3.2(C.9.3.2.12.0)	
QXCR1001241531	If a record with RDATA in excess of 65535 bytes is loaded into a nameserver, a subsequent query for that record will cause named(1M) to exit. It was a design limitation in named(1M) for handling a query in such scenario.
QXCR1001241535	named(1M) hangs, not responding to queries or control commands. Specific combinations of RDATA are loaded through cache/authoritative zone to named(1M) and a subsequent query is made.
Defects fixed in BIND 9.3.2 (C.9.3.2.11.0)	
QXCR1001225472	Incorrect data handling causes named to terminate unexpectedly

Table 9 Defects Fixed in the HP-UX in HP-UX 11i v1 and HP-UX 11i v2 operating systems *(continued)*

QXCR1001204866	Reverse lookup of the IPv4 DNS server, by executing nslookup(1), results in failure, in case pointer record for the DNS server IP address is missing from the query database, which in turn results in the failure of resolving of DNS server
QXCR1001075534	Reverse lookup of the IPv6 DNS server, by executing nslookup(1), results in failure, in case pointer record for the DNS server IP address is missing from the query database, which in turn results in the failure of resolving of DNS server
Defects fixed in BIND 9.3.2 (C.9.3.2.10.0)	
QXCR1001180790	The named(1M) crashes, if subsequent queries are sent to cache an invalid record.
Defects fixed in BIND 9.3.2 (C.9.3.2.9.0)	
QXCR1001118561	When multiple threads try to acquire locks, there is a race, because the lock handling mechanism is not synchronized. This causes the named daemon to abort.
QXCR1001154168	The named daemon allows specific packets to servers that causes named(1M) to crash
QXCR1001134927	A 64-bit binary for named(1M) is not available with BIND 9.3.2 on HP-UX 11i v2.
Defects fixed in BIND 9.3.2 (C.9.3.2.8.0)	
QXCR1001090076	When the rndc utility is terminated, named(1M) does not shut down.
QXCR1001079458	While reloading the server using rndc(1), named(1M) consumes more memory.
QXCR1001092088	While using recursive resolvers with the DNSSEC validation option enabled, the DNSSEC query to the nameserver returns a message of SERVFAIL or validation failure.
QXCR1001092086	When old signatures are retained in the named cache, validation failures occur.
Defects fixed in BIND 9.3.2 (C.9.3.2.7.0)	
QXCR1000952300	The named daemon does not behave as expected for certain messages.
QXCR1000991848	The nameserver caches invalid responses from the additional section of the response packet while processing recursive client queries.
QXCR1001009615	The nameserver sometimes returns invalid CNAME or DNAME responses.
QXCR1001004094	The nameserver sometime returns invalid NXDOMAIN responses.
QXCR1000962881	When named is started in a chroot environment, the following error is displayed: `open(/dev/poll) failed:No such file or directory`
Defects fixed in BIND 9.3.2 (C.9.3.2.4.0)	
QXCR1000848700	Some DNS responses arriving at the host are not being delivered to the `/usr/sbin/ named` process but instead are directed to other processes running on the same host.

Table 9 Defects Fixed in the HP-UX in HP-UX 11i v1 and HP-UX 11i v2 operating systems *(continued)*

QXCR1000879111	The TCP `accept()` call fails to create the new connection socket and logs one of the following errors
	`internal accept: accept() failed: Too many open files`
	`internal_accept: fcntl() failed: Too many open files`

QXCR1000848714	The closure criteria for sockets lead to inconsistent states in the socket module.
QXCR1000886576	Using the `rrset-order` option with value `fixed` in the `/etc/named.conf` file displays the following error message: `rrset-order: order 'fixed' not fully implemented.`
QXCR1000893386	The return values from the OpenSSL library functions are not checked properly in DNS code.
QXCR1000874093	Some filesets are not installed when BIND 9.3.2 is upgraded to a higher version
QXCR1000924015	DNSSEC Lookaside Validation (DLV) processing does not handle unknown signature algorithms correctly.
Defects fixed in BIND 9.3.2 (C.9.3.2.3.0)	
QXCR1000577501	The `rndc` recursing output file named.recursing contains old data.
QXCR1000821672	Forgery resilience needs more improvements.
Defects fixed in BIND 9.3.2 (C.9.3.2.2.0)	
JAGag45362	Query ID generation is cryptographically weak.
Defects fixed in BIND 9.3.2 (C.9.3.2.1.0)	
JAGag32951	`named`(1M) does not handle queries of type ANY properly.
JAGag32950	`named`(1M) aborts unexpectedly under certain circumstances.
JAGag24093	Under certain circumstances, DNSSEC utilities do not work properly.
Defects fixed in BIND 9.3.2 (C.9.3.2.0.0)	
JAGag14593	BIND does not handle SIG records properly.
JAGag14592	BIND does not handle recursive queries properly.
JAGag07595	BIND 9.x does not handle AXFR/IXFR responses properly in certain scenarios.
JAGaf71605	BIND 9.3.2 must be enabled on the HP-UX 11i v1 and v2 operating systems.

Defects fixed in the HP-UX 11i v3 operating system

Table 10 lists the defects fixed in BIND 9.3.2 in the HP-UX 11i v3 operating system.

Table 10 Defects Fixed in the HP-UX 11i v3 Operating System

Identifier	Description
Defects fixed in BIND 9.3.2(C.9.3.2.15.0)	

Table 10 Defects Fixed in the HP-UX 11i v3 Operating System *(continued)*

QXCR1001251264	It is a design limitation in named(1M) due to which the information can be prolonged in the cache beyond the period supposedly allowed by the TTL value, causing named(1M) to potentially return incorrect answers.
QXCR1001252460	It is a limitation in control scripts of named(1m) due to which named.64 is not getting killed.
Identifier	**Description**
Defects fixed in BIND 9.3.2(C.9.3.2.14.0)	
QXCR1001241531	If a record with RDATA in excess of 65535 bytes is loaded into a nameserver, a subsequent query for that record will cause named(1M) to exit.It was a design limitation in named(1M) for handling a query in such scenario.
QXCR1001241535	named(1M) hangs, not responding to queries or control commands.Specific combinations of RDATA are loaded through cache/authoritative zone to named(1M) and a subsequent query is made.
Defects fixed in BIND 9.3.2 (C.9.3.2.13.0)	
QXCR1001225472	Incorrect data handling causes named to terminate unexpectedly
Defects fixed in BIND 9.3.2 (C.9.3.2.12.0)	
None	No new defects are fixed in this version
Defects fixed in BIND 9.3.2 (C.9.3.2.11.0)	
QXCR1001180790	The *named*(1M) crashes, if subsequent queries are sent to cache an invalid record.
Defects fixed in BIND 9.3.2 (C.9.3.2.10.0)	
QXCR1001118561	When multiple threads try to acquire locks, there is a race, because the lock handling mechanism is not synchronized. This causes the `named` daemon to abort.
QXCR1001154168	The `named` daemon allows specific packets to servers that causes *named*(1M) to crash.
Defects fixed in BIND 9.3.2 (C.9.3.2.9.0)	
QXCR1001090076	When the `rndc` utility is terminated, *named*(1M) does not shut down.
QXCR1001079458	While reloading the server using *rndc*(1), *named*(1M) consumes more memory.
QXCR1001092088	While using recursive resolvers with the `DNSSEC validation` option enabled, the DNSSEC query to the nameserver returns a message of SERVFAIL or validation failure.
QXCR1001092086	When old signatures are retained in the *named* cache, validation failures occur.
Defects fixed in BIND 9.3.2 (C.9.3.2.8.0)	
QXCR1000952300	The named daemon does not behave as expected for certain messages.
QXCR1000991848	The nameserver caches invalid responses from the additional section of the response packet while processing recursive client queries.
QXCR1001009615	The nameserver sometimes returns invalid CNAME or DNAME responses.
QXCR1001004094	The nameserver sometime returns invalid NXDOMAIN responses.
QXCR1000962881	When named is started in a chroot environment, the following error is displayed: `open(/dev/poll) failed: No such file or directory`
Defects fixed in BIND 9.3.2 (C.9.3.2.6.0)	
QXCR1000848700	Some DNS responses arriving at the host are not being delivered to the `/usr/sbin/named` process but instead are directed to other processes running on the same host.

Table 10 Defects Fixed in the HP-UX 11i v3 Operating System *(continued)*

QXCR1000879111	The TCP `accept()` call fails to create the new connection socket and logs one of the following errors:
	`internal accept: accept() failed: Too many open files`
	`internal_accept: fcntl() failed: Too many open files`

QXCR1000848714	The closure criteria for sockets lead to inconsistent states in the socket module.
QXCR1000886576	Using the `rrset-order` option with value `fixed` in the `/etc/named.conf` file displays the following error message: `rrset-order: order 'fixed' not fully implemented.`
QXCR1000893386	The return values from the OpenSSL library functions are not checked properly in DNS code.
QXCR1000924015	DNSSEC Lookaside Validation (DLV) processing does not handle unknown signature algorithms correctly.
Defect fixed in BIND 9.3.2 (C.9.3.2.4.0 and C.9.3.2.5.0)	
QXCR1000841386	The local IPv6 unicast addresses, such as `fd00::/7`, are forwarded to the root server for resolution.
Defects fixed in BIND 9.3.2 (C.9.3.2.3.0)	
QXCR1000821672	Forgery resilience needs more improvements.
QXCR1000577501	The *rndc*(1) recursing output file `named.recursing` contains old data.
QXCR1000791343	*named*(1M) fails with an out of memory error message.
Defects fixed in BIND 9.3.2 (C.9.3.2.2.0)	
The defects fixed in the C.9.3.2.2.0 version are the same as the defects fixed in the C.9.3.2.1.0 version of BIND 9.3.2.	
Defects fixed in BIND 9.3.2 (C.9.3.2.1.0)	
JAGag41036	*named*(1M) fails with an "out of memory" error message if the size of the cache memory exceeds 1 GB.
JAGag45362	Query ID generation is cryptographically weak.
JAGag32951	*named*(1M) does not handle queries of type `ANY` properly.
JAGag32950	*named*(1M) unexpectedly aborts under certain circumstances.