HITACHI
Inspire the Next

HNAS Multi-Tenancy Implementation and Best Practice Guide

By Gary Mirfield

Hitachi Data Systems

MK-92HNAS059-01

Notices and Disclaimer

Hitachi Data Systems products and services can be ordered only under the terms and conditions of Hitachi Data Systems' applicable agreements. The use of Hitachi Data Systems products is governed by the terms of your agreements with Hitachi Data Systems.

This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (http://www.openssl.org/). Some parts of ADC use open source code from Network Appliance, Inc. and Traakan, Inc.

Notice of Export Controls

Export of technical data contained in this document may require an export license from the United States government and/or the government of Japan. Contact the Hitachi Data Systems Legal Department for any export compliance questions.

Document Revision Level

Revision	Date	Description
MK-92HNAS059-00	December 2014	First publication
MK-92HNAS059-01	February 2015	Second publication

Contributors

The information included in this document represents the expertise, feedback, and suggestions of a number of skilled practitioners. The author would like to recognize and sincerely thank the following contributors and reviewers of this document

- Bent Knudsen
- Bertrand Le Quellec
- Nathan King
- Phil Wood
- Troy Pillon
- Victor Abyad
- Andy Chittenden
- Ashwin Payyanadan

Contact

Hitachi Data Systems
2845 Lafayette Street
Santa Clara, California 95050-2627
https://portal.hds.com

North America: 1-800-446-0744

Table of Contents

Introduction

The Hitachi NAS Platform (HNAS) is a versatile, intelligent, and scalable multi-protocol solution. HNAS can be configured in many ways to meet a variety of requirements. This document details the best practices for configuring and using HNAS Multi-Tenancy and related features, and EVS security. It is assumed that you are familiar with networking concepts and the HNAS platform.

HNAS EVS concepts

HNAS platform-based solutions use “Enterprise Virtual Servers” (EVSs) to provide file services to hosts. A server or cluster supports up to 64 EVSs. Each EVS is assigned unique network settings and storage resources, enabling administrators the flexibility to logically partition access to shared storage resources. In clustered HNAS systems, EVSs are automatically migrated between servers when faults occur to ensure maximum availability.

By deploying multiple EVSs on one HNAS appliance, customers have the ability to provide secure multi-tenancy solutions for complex environments supporting separate networks and security contexts. While running on the same hardware all services are completely separated and no internal communication between EVSs is possible.

The Hitachi Data Systems High Performance NAS appliance can be configured in a number of different configurations including:

- Separate non-overlapping tagged VLANs on the same network from a single HNAS node or cluster with EVS authentication to a single AD/LDAP entity.
- Separate physical networks from a single HNAS node or cluster using multiple port aggregations to a single AD/LDAP entity.
- Separate overlapping tagged VLANs on the same network with per-EVS authentication to multiple AD/LDAP entities using the Virtual Server Security License.
- Separate physical networks from a single HNAS node or cluster using multiple port aggregations with per EVS authentication to multiple AD/LDAP entities using the Virtual Server Security License.

The Virtual Server Security License enables the creation of Virtual Servers in an individual security context, allowing each HNAS Virtual Server to authenticate to separate Active Directories and/or NIS-LDAP domains. This is an essential feature for ISPs offering hosting or outsourcing services where multiple end customers are sharing the infrastructure, or for customers wanting to separate their Production and Test/Development facilities. Administrators may specify distinct resources for each Virtual Server and allows the Hitachi HNAS Platform to participate in up to 64 independent non-trusted Windows Active Directory Top-level Domains or up to 64 independent NIS domains unless distinct resources have to be defined for each Virtual Server where no trusts exist.

Multi-Tenancy

Overview

Multi-tenancy is a configuration mode used together with the EVS Security License to support multiple tenants, where each EVS may use unique security and network contexts within the same physical server or cluster. Multi-tenancy extends HNAS individual security mode to provide true separation by maintaining per-EVS variables and connection states. Serving environments for

tenants can be single or multiple EVSs configured separately possibly sharing file serving interfaces. This configuration mode enables support for overlapping IP addressing.

HNAS multi-tenancy is intended for the following use cases:

- Application Service Providers (ASPs) providing NAS services to Cloud consumers.
- Customers with complex environments comprising of multiple networks and authentication methods
- Customers with Production, QA and Dev environments requiring segmentation and additional security between environments

Figure 1: Multi-tenancy configuration

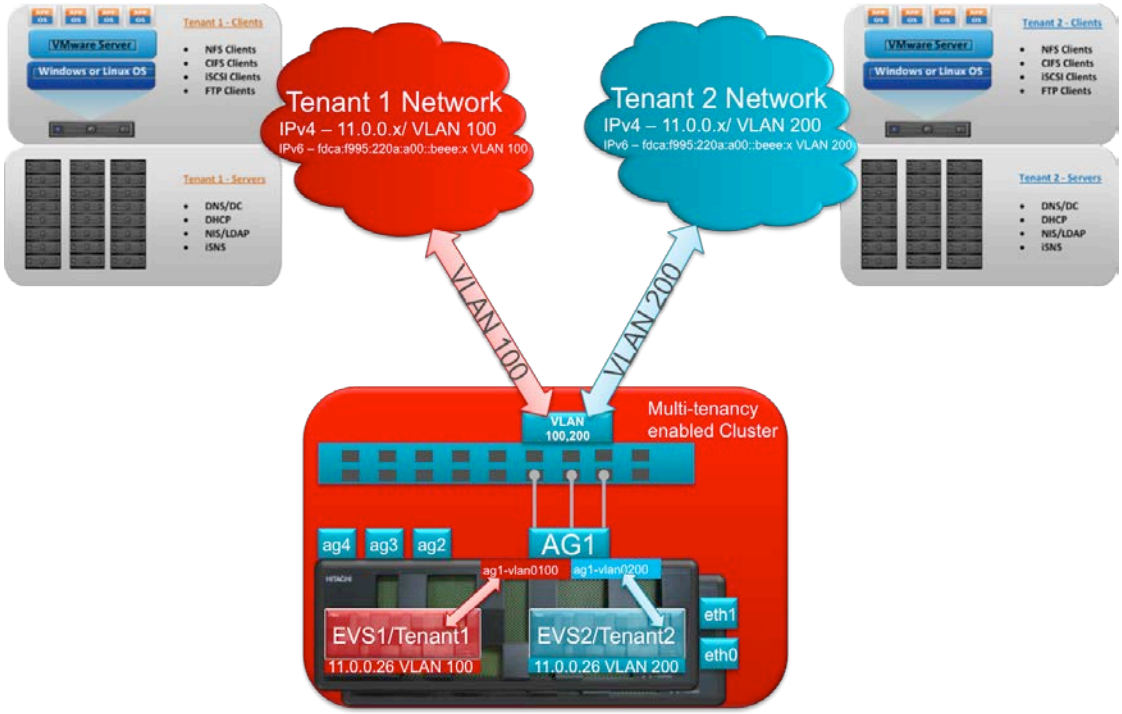

Incoming requests and outgoing responses are made with the EVS in context. The solution provides security between environments by performing crosschecks of server addressing, interfaces and virtual nodes as data is passed between stacks.

The following protocols are supported for use with multi-tenancy:

- SMB/CIFS
- NFS
- iSCSI
- FTP

Note: The following features are not supported for use in a multi-tenancy environment:

- Data Migrator to Cloud

- Cluster Name Space (EVS namespaces are permitted)
- File System Relocation (migration between EVSs)
- NetBIOS
- SSC connectivity to File Serving IP addresses

Data Migrator to Cloud does not support the use of VLAN interfaces in release 12.3.

The Cluster Namespace feature is designed to offer a cluster wide virtual namespace of resources across multiple EVSs within a cluster. As multi-tenancy implementations are designed to use EVSs in separate security contexts, the feature has limited value for these use cases.

Due to the lack of NetBIOS support, clients running operating systems earlier than Windows 2000 and OS X Lion could be affected.

All Generic Management Server (GMS) and SSC connections to File Serving IP addresses are disabled when multi-tenancy is enabled. This is to prevent a tenant from becoming an administrator on the HNAS and gaining access to other tenants' configuration and data.

HNAS multi-tenancy is fully compatible with other HNAS features, such as Snapshots and File- or Object-based Replication.

HNAS Networking

Link Aggregation

HNAS uses the term Link Aggregation (LAG) to describe combined physical network ports combined for the purposes of load balancing and redundancy. Depending on the environment into which the HNAS solution will be integrated, static or dynamic Link Aggregation may be used.

Static Links

- Sufficient for Single Core Switches
- A grouping of ports specifically configured to load share.
- The switch ports at each end must be configured as part of a load-sharing group.
- Additionally, you can choose the load-sharing algorithm used by the group.

Link Aggregation Control Protocol (LACP)

- Single Logical Link (802.3ad) between the Switch and HNAS
- HNAS controls the load balancing I/O
- Packet traffic is distributed among the links using a hash function based on a variety of parameters, including source and destination MAC addresses.

HNAS link aggregations can also support certain types of Multi-Chassis Link Aggregation Groups. These technologies are proprietary to individual switch vendors, such as Virtual Link Trunking (Dell), Virtual Port Channel (Cisco), DNOS6.x (Dell) or Intelligent Resilient Framework (HP). For more information, refer to vendor specific documentation and MK-92USF003, "Hitachi Unified Storage File Module Network Administration Guide."

VLAN Tagging

HNAS optionally may use VLAN tagging to minimize the number of switches, ports/NICs required in an environment. VLAN tagging allows VLAN aware equipment to share physical ports while maintaining separation of the broadcast domains. VLANs typically map to IP Subnets but not always; a VLAN may contain multiple subnets. Ethernet frames have a tag applied to identify their VLAN, which is then stripped upon exiting the VLAN.

Figure 2: HNAS with switch port default VLAN Assignment

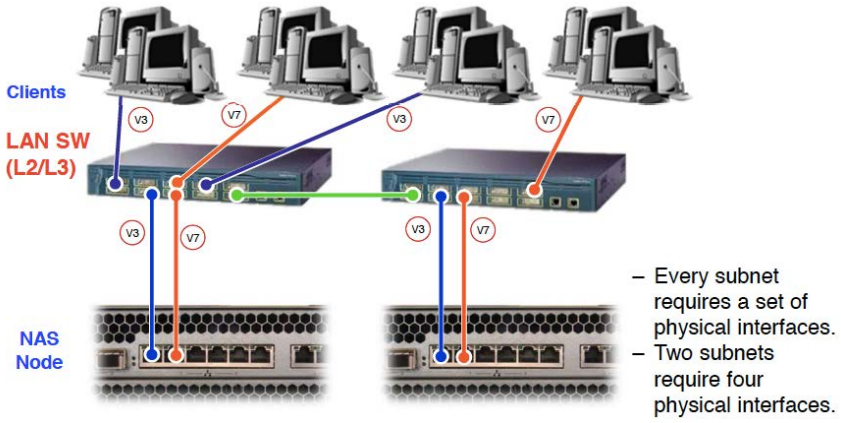

In this example, each separate subnet requires a minimum of one dedicated interface placed into a unique link aggregation group. Therefore, two subnets require four physical interfaces.

Figure 3: HNAS with switch port VLAN tagging enabled

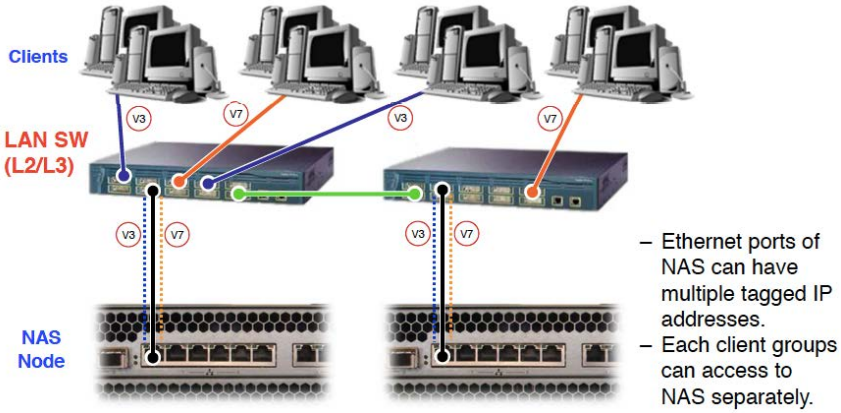

In this example, multiple tagged VLANs may share the same interfaces and link aggregation groups. It is also possible for both tagged and untagged (default) VLANs to be used simultaneously by the same switch ports.

Interface Groups

In HNAS v12.0 and higher, the concept of interface groups has been extended from link aggregation groups to include a specific VLAN description. This also introduced the concept of per-EVS routing, working along with the per-node and cluster wide routing capabilities available in

earlier code releases. While enabling per-EVS routing does not require the EVS security license, overlapping subnets are not permitted without enabling multi-tenancy.

Figure 4: Interface format

Overlapping Subnets

Enabling multi-tenancy will also enable support for configuring duplicate/overlapping server IP addresses/ IP subnets governed by the following rules:

- The same IP address is not used by the same EVS more than once.
- The same IP address is not used by the same interface more than once where an interface can be an aggregate or a VLAN Interface.

Figure 5: Overlapping Subnets

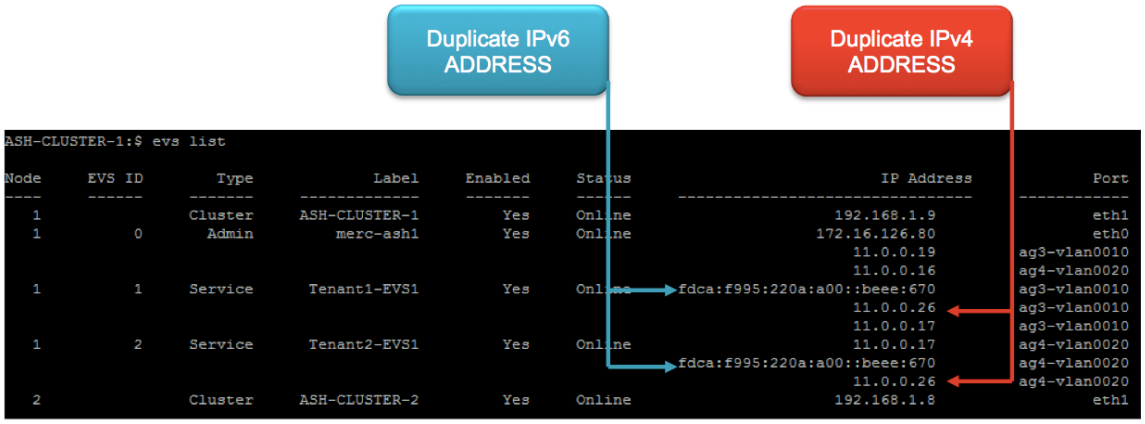

Implementation

Licensing

HNAS multi-tenancy is an extension of EVS individual security mode and installing the EVS security license makes the HNAS system capable of multi-tenancy. Check the list of licensed features on the HNAS system used to ensure the correct license is present.

Figure 6: HNAS Licences

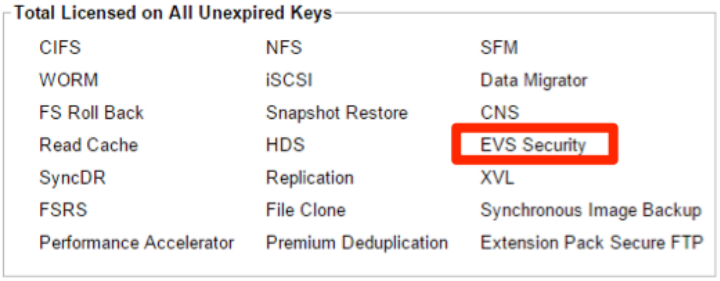

Switch Configuration

Before configuring VLAN interfaces on the HNAS, the switch configuration should be verified. The `show vlan` command may be used on Cisco Nexus switches to display a list of configured VLANs and their permitted ports. In this example, VLANs 100 and 149 are permitted as tagged VLANs, while VLAN 161 is assigned to the same port as untagged.

```
s55-07#show vlan

Codes: * - Default VLAN, G - GVRP VLANs, P - Primary, C - Community, I -
Isolated

Q: U - Untagged, T - Tagged

   x - Dot1x untagged, X - Dot1x tagged

   G - GVRP tagged, M - Vlan-stack, H - VSN tagged

   i - Internal untagged, I - Internal tagged

    NUM         Status    Description      Q Ports

    100         Active                     T Te 0/48

    149         Active                     T Te 0/48

    161         Active                     U Te 0/48
```

VLAN Interfaces

Multi-tenancy is supported from HNAS release 12.2 onward. For systems upgraded from 11.x or earlier, multi-tenancy requires any VLAN configuration created using the deprecated `vlan` information to be updated to use the new `vlan-interface` command set. The `vlan-interface-*` suite of commands is used to create and modify VLANs in HNAS release 12.0 and greater.

```
HNAS4100-1:$ $ vlan-interface-create --interface ag1 161

Created ag1-vlan0161
```

The `vlan-convert-config.rb` script may also be used to generate the HNAS CLI command sequence required to convert an existing VLAN-subnet configuration. The script does not run under the HNAS CLI but is available on the HNAS platform, from HNAS Version 12.2, through the Linux console and is located at: `/opt/mercury-utils/bin/vlan-convert-config.rb`

After reconfiguration, the command `vlan-interface-show` will display an updated format of <aggregate-name>-<vlan-id>, as shown below.

```
HNAS4100-1:$ $ vlan-interface-show
ag1-vlan0161
```

EVS Configuration and Security

Where necessary, create all necessary File Serving EVSs using the `evs` command. The –p parameter specifies the VLAN interface.

```
HNAS4100-1:$ $ evs create –l EVS4 –i 172.17.161.23 –m 255.255.255.0 –p ag1-vlan0161
EVS 4 created successfully.
```

EVS individual security mode is a prerequisite (for all configured EVSs) for enabling multi-tenancy. To convert existing EVSs from a global to individual mode, first disable each EVS using the `evs-disable` command, then convert using the `evs-security individual` command and re-enabling the EVS using `evs enable`. Repeat this process for all EVSs.

```
HNAS4100-1:$ evs disable -e 1
Do you want to proceed?(Y/N)[N]:
Y
HNAS4100-1:$ evs-security individual -e 1
HNAS4100-1:$ evs enable -e 1
```

Routing

Most multi-tenancy use cases will also require the `routing-by-evs` capability to be enabled. Enabling the multi-tenancy feature will also automatically enable `routing-by-evs.` Once enabled, the choice of source addresses available to the routing engine is restricted to those associated with that EVS. To manually enable `routing-by-evs` prior to multi-tenancy, the routing`-by-evs-enable` command may be used.

```
HNAS4100-1:$ routing-by-evs-enable
HNAS4100-1:$ routing-by-evs-show
routing-by-EVS is enabled
```

Gateway, network and host routes (IPv4 and IPv6) are configured per EVS once `routing-by-evs` is enabled using the following commands:

- `route-gateway-add`
- `route-net-add`
- `route-host-add`

Routing should be configured by using the cli in an EVS context. For example:

```
HNAS4100-1:$ evs-select 2
HNAS4100-1[EVS1]:$ route-net-add 10.2.0.0/16 -g 10.1.2.3 -m 9000
```

or

```
HNAS4100-1:$ vn 2 route-net-add 10.2.0.0/16 -g 10.1.2.3 -m 9000
```

Routes may be added for both IPv4 and IPv6, at default gateway, network or host level as before.

```
HNAS4100-1:$ vn 2 route-net-add 10.2.0.0/16 -g 10.1.2.3 -m 9000
HNAS4100-1:$ vn 1 route-gateway-add fdca:f995:220a:a00::1
HNAS4100-1:$ vn 3 route-host-add 10.1.2.3 -g 10.1.3.4
```

The `route` command is for display purposes only once multi-tenancy is enabled and will display routes for the EVS in context.

```
HNAS4100-1:$ evs-select 2
HNAS4100-1[EVS1]:$ route
Routes for EVS 2:
Destination                    Gateway                    MTU     Flags
0.0.0.0/0                      172.16.1.1                 Default G
::/0                           fe80::208:e3ff:feff:fc28   1500    GD
via eth0
::/0                           fe80::208:e3ff:feff:fc28   1500    GD
via ag2
::/0                           fe80::208:e3ff:feff:fc28   1500    GD
via ag1
```

Attempting to use the route command for configuration will redirect the user to the route-* commands for configuration when multi-tenancy is enabled.

```
HNAS4100-1[EVS1]:$ route add gateway
route: as multi-tenancy is enabled, use route-gateway-add
```

Per-EVS routes are limited to a total of 127 static or MTU routes.

Enable Multi-Tenancy

To enable multi-tenancy on a new system where no EVSs are present, or EVSs are already in individual security contexts the command `multi-tenancy-enable` may be used.

```
HNAS4100-1:$ multi-tenancy-enable
Warning: Enabling multi-tenancy significantly affects the configuration of
the HNAS.
Have you read and understood the multi-tenancy man page?(Y/N)[N]:
Y
Have you read and understood the multi-tenancy-enable man page?(Y/N)[N]:
Y
Do you understand that once enabled, multi-tenancy cannot be disabled until
all file serving EVSs have been deleted?(Y/N]:
Y
Warning: All active connections, including any remote console sessions,
will be disconnected to allow the network service to support multi-tenancy.
Do you want to proceed?(Y/N)[N]:
Y
Enabling multi-tenancy.
Operation successful.
[multi-tenancy-enable took 16 s.]
HNAS4100-1:$ Connection closed by foreign host.
```

Note: Enabling multi-tenancy will cause a temporary loss of service when the feature is enabled.

The status of multi-tenancy can be verified using the `multi-tenancy-show` command and if necessary disabled using `multi-tenancy-disable`.

Note: Once enabled, multi-tenancy cannot be disabled until all file serving EVSs have been deleted.

Manageability

An improvement made in the 12.1 HNAS code release provided a function within the HNAS cli prompt to display the EVS console context within square brackets for ease of use. To set the EVS name in context on the CLI, execute `evs-select <EVS ID>` to run all commands in context, or alternatively use vn `<EVS ID> <command>` on a per command basis.

```
HNAS4100-1:$ evs-select 2
HNAS4100-1[EVS1]:$
```

DNS

After enabling multi-tenancy, use the CLI to configure DNS for each EVS using the `dnsserver` and `dnsdomainname` commands.

```
HNAS4100-1:$ vn 1 dnsserver add 172.31.60.10 172.31.60.11
HNAS4100-1:$ vn 1 dnsdomainname set tenant1.local
HNAS4100-1:$ vn 2 dnsserver add 172.31.60.10 172.31.60.11
HNAS4100-1:$ vn 2 dnsdomainname set tenant2.local

HNAS4100-1[EVS1]:$
```

Active Directory

Each EVS may now be joined into a unique Active Directory domain. This may be performed using the Add CIFS Server Names page of the SMU GUI.

Figure 7: Active Directory configuration

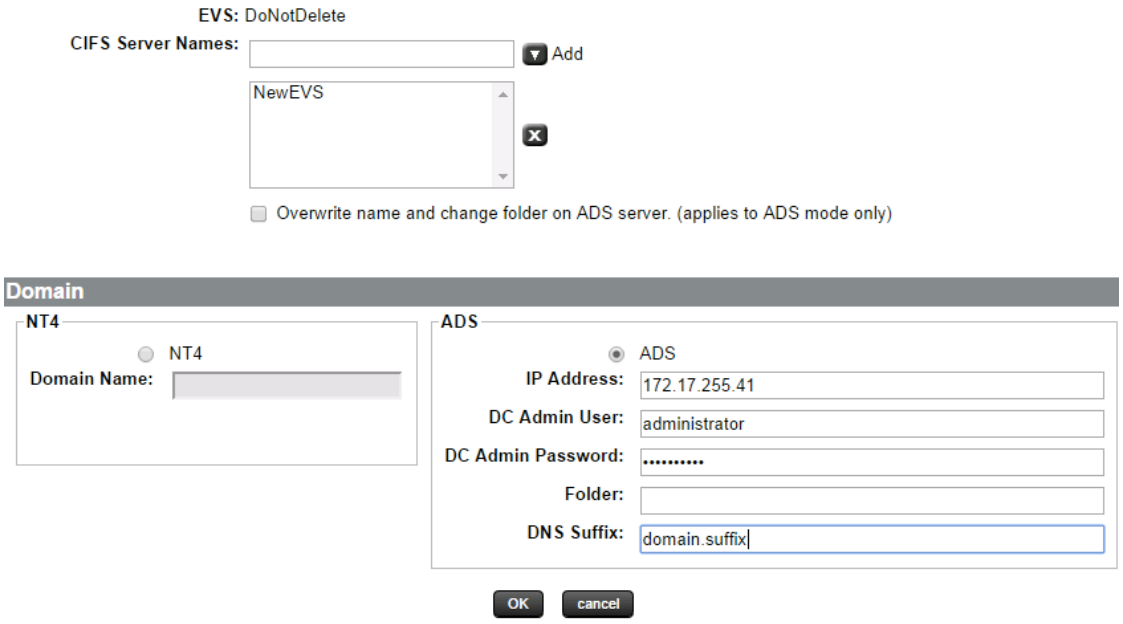

Best Practices

Best practices are identified throughout the document. This section discusses best practices not covered elsewhere.

EVS Naming Conventions

In some environments, it may be desirable to provision multiple EVSs from a single cluster which share resources such as network VLANs, Active Directory and DNS resources. Prior to Multi-Tenancy, a single group of these EVSs may have utilized the Global Security context. With Multi-Tenancy enabled, every EVS uses a separate security context but may participate in common VLANs. From an ease of use perspective, consider implementing an EVS naming convention which reflects this in order to benefit from the EVS context display introduced as a CLI enhancement. An example such as <End Customer>-<CIFS Serving Name> could be effective.

EVS Design

Do not create more EVSs than necessary to support a particular environment and observe existing platform limitations for concurrent user sessions. Concurrent user sessions are calculated on a per-EVS basis. For example, a 20k user CIFS environment with a single EVS may generate up to 20k concurrent connections. If those 20k users are using shared resources from two EVSs concurrently, the total connection count for that tenant could theoretically reach a maximum of 40k.

Heap Memory Allocation

Use the `resource-config-*` commands to list available heap memory during implementation and optimize the configuration only if necessary. For environments which are not using features such as de-duplication or data migrator to cloud extensively, the heap allocation may be increased to support additional user connections and open files.

Model	Default Configuration		Adjusted Heap Configuration	
	# connections	# open files	# connections	# open files
4060	45 000	227 000	64 000	565 000
4080	45 000	227 000	64 000	565 000
4100	64 000	1 178 000	64 000	1 500 000

```
HNAS4100-1:$ resource-config-show
Current node 1:
                Memory      Current
----------------------     --------
       System overhead     1.598 GB
             Bali heap        24 GB
         Deduplication       [auto]
Data Migrator to Cloud       [auto]
          Embedded SMU       [auto]
       Total available        32 GB
```

```
       Total committed      25.6 GB
    Uncommitted memory     6.402 GB


HNAS4100-1:$ resource-config-maximum-safe-heap-show
Maximum safe Bali heap: 31656120320 B (29.48 GB)
Parameter group for use with resource-config commands: "-f bali-heap
31656120320"


HNAS4100-1:$ resource-config-apply -f bali-heap 31656120320
Current node 1:
                Memory     Previous      Current     Restarted
----------------------     --------     --------     ---------
       System overhead     1.598 GB     1.598 GB      1.598 GB
             Bali heap        24 GB        24 GB      29.48 GB
          Deduplication       [auto]       [auto]        [auto]
Data Migrator to Cloud       [auto]       [auto]        [auto]
           Embedded SMU       [auto]       [auto]        [auto]
       Total available        32 GB        32 GB         32 GB
       Total committed      25.6 GB      25.6 GB      31.08 GB
    Uncommitted memory     6.402 GB     6.402 GB      942.4 MB

Restart is required on all nodes to fully apply memory settings.
HNAS4100-1:$
```

Hitachi Data Systems

Corporate Headquarters
2845 Lafayette Street
Santa Clara, California 95050-2639
U.S.A.
www.hds.com

Regional Contact Information

Americas
+1 408 970 1000
info@hds.com

Europe, Middle East, and Africa
+44 (0) 1753 618000
info.emea@hds.com

Asia Pacific
+852 3189 7900
hds.marketing.apac@hds.com

MK-92HNAS059-01