

CUDB Node Network Description

USER GUIDE

Copyright

© Ericsson AB 2016-2018. All rights reserved. No part of this document may be reproduced in any form without the written permission of the copyright owner.

Disclaimer

The contents of this document are subject to revision without notice due to continued progress in methodology, design and manufacturing. Ericsson shall have no liability for any error or damage of any kind resulting from the use of this document.

Trademark List

All trademarks mentioned herein are the property of their respective owners. These are shown in the document Trademark Information.



Contents

1	Introduction	1
1.1	Document Purpose and Scope	1
1.2	Revision Information	1
1.3	Target Groups	3
1.4	Typographic Conventions	3
2	Solution Overview	5
2.1	CUDB System Overview	5
2.2	Network Overview	5
2.3	VIP Address Allocation	26
2.4	Address Plan Summary	29
2.5	Time Synchronization	32
2.6	Traffic Flows	32
2.7	Routing	40
2.8	Firewall Configuration	41
3	Appendix: Quality of Service	43
	Glossary	45
	Reference List	47





1 Introduction

This document describes the network details of an Ericsson Centralized User Database (CUDB) node, and gives an overview of the interwork among the CUDB nodes comprising the CUDB system.

1.1 Document Purpose and Scope

The purpose of this document is to describe the general network infrastructure of a CUDB system, from the internal components to the node integration with the customer network environment. Most of the network details covered by this document are independent of the available CUDB infrastructure types.

The document covers the following topics in detail:

- Overview of the CUDB node and CUDB system network infrastructure.
- The logical networks of the CUDB node, both internal and external.
- Logical connectivity.
- IP addressing plan.

The following topics are, however, out of the scope of this document.

- Specific network configuration details.
- Details of the customer network outside the CUDB node.
- Customer-tailored network integration for system management or backups. This is part of the customer integration project.

1.2 Revision Information

Rev. A

This document is based on 4/1551-CSH 109 067/9 with the following changes:

- Document revised and expanded.
- Removed obsolete information.
- Virtualization terminology update throughout the document.

Rev. B

Editorial changes only.



Rev. C

Other than editorial changes, this document has been revised as follows:

- Section 2.2.3.2.1 on page 15 and Section 2.4 on page 29: Added information regarding customer specific deployments.
- Section 2.6.1.1 on page 33: Updated Table 3, OAM NETCONF traffic type is enabled on SITE network.
- Section 3 on page 43: Updated Table 18, SITE_VIP network is added to CUDB_OAM In NETCONF traffic flow.

Rev. D

Other than editorial changes, this document has been revised as follows:

- Section 2.2.3.3 on page 19: Added information on CUDB systems deployed on native BSP 8100 with GEP3 blades.
- Section 2.2.3.4 on page 22: Added information on configuring the redundancy mechanism.

Rev. E

Other than editorial changes, this document has been revised as follows:

- Section 2.6 on page 32: Updated description.
- Section 2.6.1 on page 33: Updated header in Table 3, Table 4, Table 5, Table 6, and Table 7.
- Section 2.6.2 on page 36: Updated description.
- Section 2.6.2.1 on page 36: Added information about authentication LDAP queries to Table 8.
- **Section 2.6.2.9 CUDB OAM Centralized Authentication System Support Traffic:** Removed section.
- Section 2.8 on page 41: Updated LDAP/LDAPS traffic networks in Table 18.

Rev. F

Other than editorial changes, this document has been revised as follows:

- Section 2.2.2.2 on page 9: Updated Front-End Element and Gateway Routers terms.
- Section 2.2.3.2.2 on page 17: Updated description.
- Section 2.6.2.9 on page 40: Added section with Table 16.



- Section 2.7 on page 40: Updated routing configuration constraints.

Rev. G

Other than editorial changes, this document has been revised as follows:

- Section 2.3.1.1 on page 27: Updated description and step list.

Rev. H

Other than editorial changes, this document has been revised as follows:

- Section 2.4 on page 29: Updated vCUDB_FE, vCUDB_OAM, vCUDB_PROVISIONING, and vCUDB_SITE Net/Mask examples in Table 2 and added note.

Rev. J

Other than editorial changes, this document has been revised as follows:

- Section 2.4 on page 29: Added table notes and updated the Purpose of SysMGMT Network Common Name in Table 1, and updated IPv6 examples of vCUDB_FE, vCUDB_OAM, vCUDB_PROVISIONING, and vCUDB_SITE Network Common Names in Table 2. Added information about IPv6 Support.

1.3 Target Groups

This document is intended for CUDB integration personnel, system administrators, and system architects.

The document assumes a general knowledge of networking, and the knowledge of basic CUDB product architecture both at system and node levels.

1.4 Typographic Conventions

Typographic conventions can be found in the following document:

- [Typographic Conventions](#)





2 Solution Overview

This section describes the networking principles applicable to any CUDB system. Mapping these general principles to actual deployments is out of the scope of this document.

2.1 CUDB System Overview

CUDB is a distributed database system, exposed as a Lightweight Directory Access Protocol (LDAP) directory, physically made up of network-connected (and interconnected) CUDB nodes spread over the operator network. Logically and length-wise, the CUDB system is divided into two internal layers, depending on the provided function: Processing Layer (PL) and Data Store (DS) Layer.

Refer to [CUDB Technical Product Description](#), Reference [1] for more information on node components. For CUDB systems deployed on native BSP 8100, refer to [CUDB Node Hardware Description](#), Reference [2] for more information on CUDB hardware.

2.2 Network Overview

This section provides an overview of the hardware and software components used to configure the CUDB network, as well as a general network description depending on the type of deployment.

For CUDB systems deployed on native BSP 8100, the following elements are covered:

— Hardware components:

- Component Main Switch (CMX) routers, see Section 2.2.1.1 on page 6.
- System Control Switch (SCX) boards, see Section 2.2.1.2 on page 7.

— Software components:

- BSP 8100 software, see Section 2.2.2.1 on page 8.
- Evolved Virtual IP (eVIP), see Section 2.2.2.2 on page 9.

— Network:

- General overview, see Section 2.2.3.1 on page 11.
- Networks description, see Section 2.2.3.2 on page 13.
- External connectivity, see Section 2.2.3.3 on page 19.



- Bidirectional Forwarding Detection (BFD) and Virtual Router Redundancy Protocol (VRRP) protocols, see Section 2.2.3.4 on page 22.
- Internal L2 connectivity, see Section 2.2.3.5 on page 23.

— Other considerations:

- Hardware alarms, see Section 2.2.4.1 on page 26.
- Hardware naming convention, see Section 2.2.1 on page 6.

For CUDB systems deployed on a cloud infrastructure, the following elements are covered:

— Software components:

- eVIP, see Section 2.2.2.2 on page 9.

— Network:

- General overview, see Section 2.2.3.1 on page 11.
- Networks description, see Section 2.2.3.2 on page 13.

2.2.1 Hardware Components

This section describes the hardware components of the CUDB network for CUDB systems deployed on native BSP 8100.

2.2.1.1 CMX

The CMX supplies 10 Gigabit Ethernet (GbE) Switching to all the backplane ports, used for connectivity of all Generic Ericsson Processor (GEP) blades in the EGEM2 subrack, 10GbE/40GbE to four front ports, 10GbE to another four front ports, and 1GbE to four front ports, in compliance with applicable parts of the current IEEE 802.3, IEEE 802.1D and IEEE 802.1Q standards. The CMX resides on a Component Main Switch Board version 3 (CMXB3) hardware component.

The CMX software runs on the CMXB3, and provides 1GbE/10GbE/40GbE switching and routing solution for nodes based on EGEM2-10 and EGEM2-40 shelves, and Cabinet Aggregation Switch (CAX) 2.0 units in Cabinet Aggregation Shelves (CAS). In addition to L2 switching and L3 routing functions, the CMX also provides a comprehensive set of various operation and maintenance (OAM) functions, such as board management of HWM, SWM, SYS, PM, CM, SYNC, and logging.

The CMX switch is designed to run in a network environment with separated management and payload networks, or in other words, on separate control and data planes. This means that some of the interfaces are dedicated for control traffic, while others are for data traffic.



The control interfaces are divided into two categories:

- Trusted: Control interfaces towards the backplane are considered trusted, and are opened at system start by default.
- Non-trusted: Control interfaces on the front are considered non-trusted, and are closed at system start by default.



Figure 1 CMX Front

2.2.1.2 SCX

The SCX is an Ethernet switch with hardware and software functions. The backplane contains 26 1GbE ports. The front ports, at the same time, are used for the inter-System Control Switch Board (SCXB) connectivity, and connection towards the CMX.

The SCXB has seven front ports: four 10GbE ports (E1-E4) and three 1GbE ports (GE1-GE3). In the backplane, the two SCXBs of the same subrack are connected by two backplane ports for DMX monitoring. Also, the two SCXBs in the lower magazine are connected by a cable in one front port for redundancy purposes.

SCX provides functions for blade Preboot eXecution Environment (PXE) or Dynamic Host Configuration Protocol (DHCP) boot infrastructure and blade hardware equipment management in an EGEM2 subrack, such as power, fan control, supervision, and other operation and management tasks.




Figure 2 SCXB3 Front

The SCXs are internal CUDB node control plane switches only, and are not connected to external virtual networks.

Although there is an interconnection between the Ethernet bridges of the SCXBs through the backplane, it is reserved for control operations.

2.2.2 Software Components

This section lists the software components of the CUDB network.

2.2.2.1 BSP 8100 Software for CUDB Systems Deployed on Native BSP 8100

BSP 8100 is an infrastructure management product that provides the backbone for an Ericsson multi-application site solution. It offers services and resources to streamline user interface and implementation for common infrastructure tasks for different Ericsson Applications. Applications that are running on BSP are implemented as BSP Tenants.

The main objectives of BSP 8100 as a product are to provide hardware and software decoupling and also decoupling of External Network and Tenant Network Connectivity.



- Hardware and software decoupling: BSP 8100 is responsible for hardware management, that is, it is responsible for supervision and maintenance operations of the BSP hardware. The BSP Tenants run natively on the blades, that is, each BSP Tenant owns its associated blades as an execution platform.
- External Network and Tenant Network decoupling: BSP 8100 is responsible for communication towards the External Networks. External Network connectivity is logically or physically separated, depending on the operator requirements on the network infrastructure. BSP 8100 provides the connectivity to the BSP Tenants through a pair of Virtual Routers (VRs). BSP 8100 offers resources to the BSP Tenants in terms of execution platforms (device blades) and transport functionality (virtual networks, VRs, and so on). BSP 8100 also offers services in terms of Network Time Protocol (NTP), control interfaces of the BSP 8100 system (BGCI), Address Resolution Protocol (ARP) targets and more.

The BSP software consists of the following:

- DMXC SW: The DMXC is the Infrastructure Management SW of BSP 8100. It implements configuration management, fault management, equipment management, security management, and performance management for CUDB systems deployed on native BSP 8100. DMXC is a HA application using a 1+1 redundancy scheme.
- SCX SW: The SCX SW is installed on all SCX boards. The SCX SW offers the operating system for DMXC to run as well as services that allow DMXC to be installed and upgraded. The Shelf Manager (ShM) is a distributed application running on each pair of SCXs in each subrack.
- CMX SW: The CMX SW is installed on all CMX boards. It provides routing and switching functionality.
- Capturing Unit (CU) SW: The CU SW includes an IP Tracing tool that is installed on one GEP5 board and provides tracing capabilities for troubleshooting BSP Tenant traffic. The CU SW is optional and is currently not used by CUDB.

2.2.2.2

eVIP

Traffic distribution is provided by the eVIP framework. Some of the main terms related to eVIP are explained in the following list:

**Abstract Load Balancer (ALB)**

The ALB is a logical container for eVIP addresses. The ALB concept is comparable to the VR concept available in commercial routers. Load balancer functions are used to distribute traffic (for example, TCP connections) according to a distribution policy (for example, round robin) among a defined set of targets. Such a set of targets are referred to as a “pool of targets”.

The ALB also serves as a structuring entity that compartmentalizes scalable Server Load Balancing (SLB) resources and external interfaces. An ALB, therefore, can be viewed as the equivalent of a commercial SLB box which is embedded in the cluster using eVIP. However, an important difference is that commercial SLB boxes are deployed as external appliance boxes in the Customer Premises Equipment (CPE): they are individually installed as other independent network appliance boxes, such as routers and firewall boxes.

Front-End Element

The Front-End Elements handle the communication with the external Data Communication Network and are responsible for routing traffic towards and from VIP addresses.

Gateway Routers

Any component capable of doing Equal Cost Multi Path (ECMP). For CUDB systems deployed on native BSP 8100, this task is carried out by CMX routers. For CUDB systems deployed on a cloud infrastructure, Gateway Routers are external entities, not part of the Virtualized Network Function (VNF).

Load Balancer

IP Servers where load distribution is done. IP Servers are transparent, hiding all the eVIP details from applications.

Security Element

The Security Element is a part of the IPsec implementation. It applies security policies on the traffic flows, and encrypts or decrypts traffic if necessary.

As a consequence of using eVIP software, a set of additional virtual networks are required. These virtual networks are for internal purposes and not routable from the customer network.

Details about eVIP and its architecture can be found in [eVIP Management Guide](#), Reference [5].



2.2.3 Network

2.2.3.1 General Overview

For CUDB systems deployed on native BSP 8100, Figure 3 provides a general overview of how previously described components connect to each other to shape the BSP 8100 hardware network. The aim of this figure is to help understand the rest of the topics in this document, it is not an exhaustive or detailed figure.

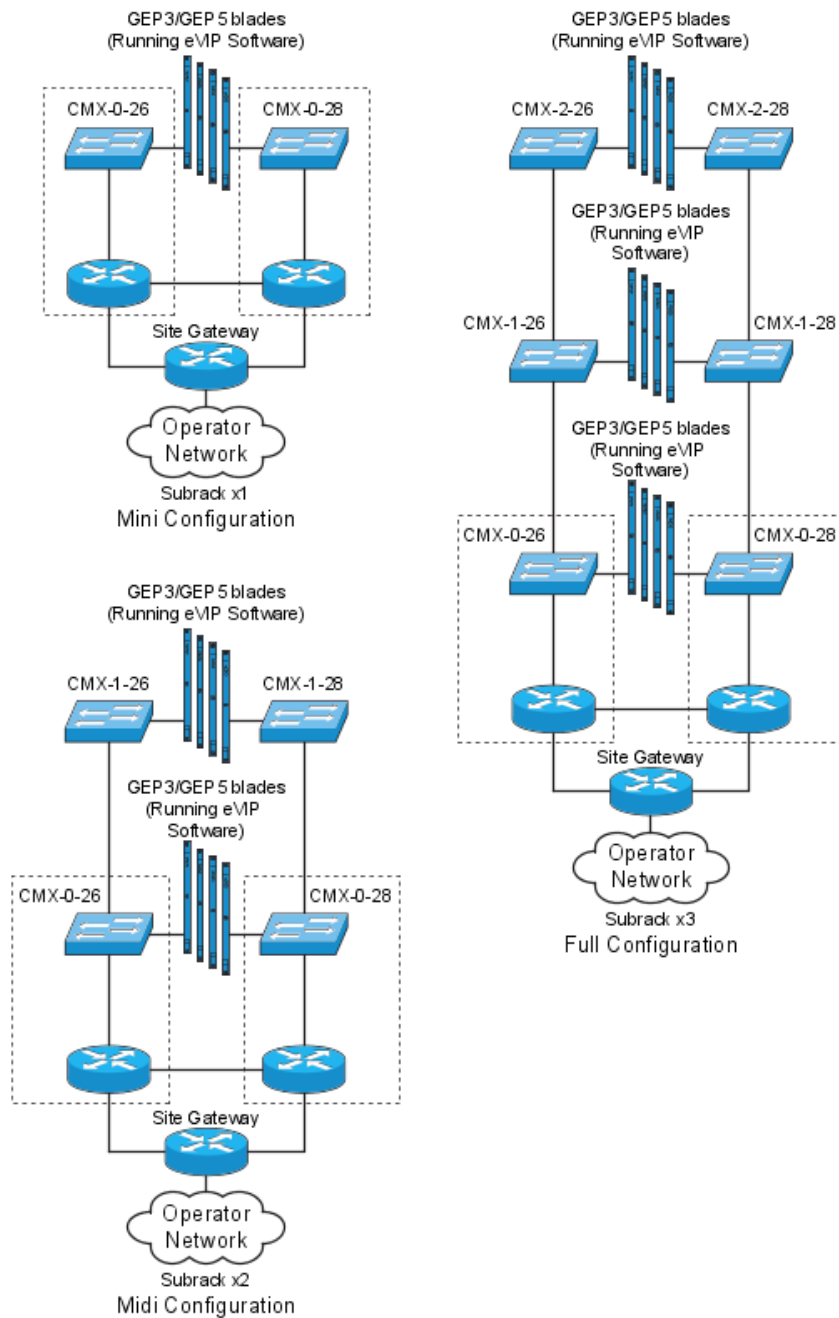


Figure 3 CUDB Node on BSP 8100 Hardware, General Network Overview

For CUDB systems deployed on a cloud infrastructure, Figure 4 provides a general overview of the virtual network layout. The aim of this figure is to help understand the rest of the topics in this document, it is not an exhaustive or detailed figure.

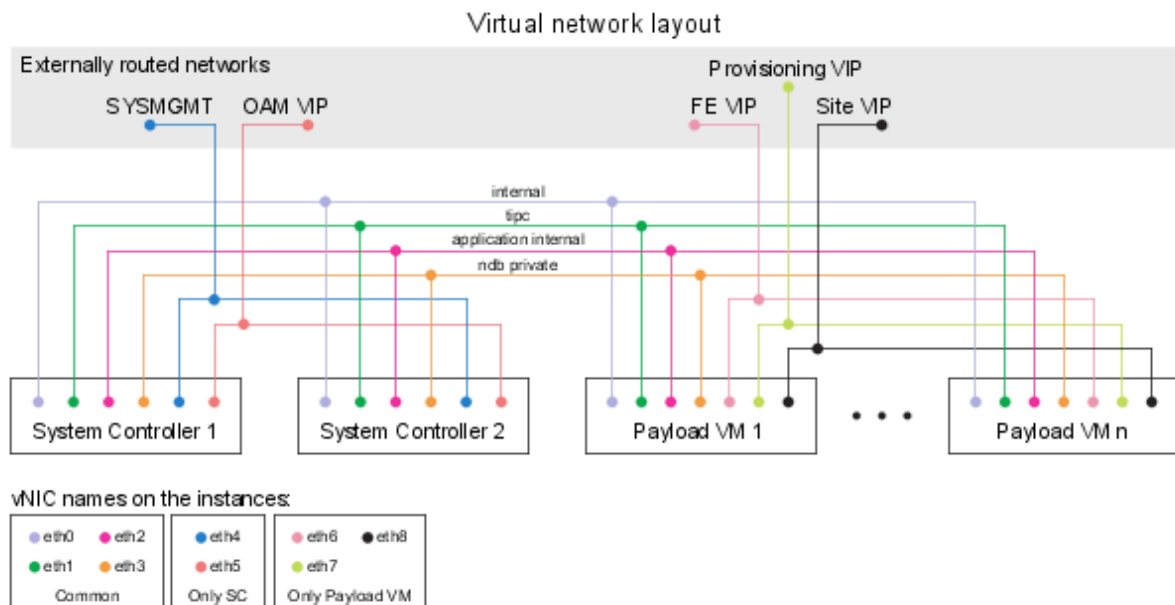


Figure 4 CUDB Node on a Cloud Infrastructure, General Network Overview

2.2.3.2 Network Description

This section gives an overview of the different internal networks configured within a CUDB node. These internal networks are used to separate different types of traffic.

Additionally, the CUDB node can be connected to several outer networks defined in the external network site infrastructure. These networks are either defined to separate OAM, provisioning, inter-CUDB, and application FE traffic, or to fulfill other reasons, like ensuring better manageability, traffic flow control, and applying QoS.

Therefore, available CUDB virtual networks are classified into the following categories, depending on their use and the region they belong to:

- **Internal:** Private networks inside the node are used to internally address the blades or Virtual Machines (VMs) in each CUDB node. Therefore, addresses within these networks are not routable outside the CUDB node. See Section 2.2.3.2.1 on page 15 for more information.
- **Infrastructure:** Infrastructure networks are used to transport incoming and outgoing CUDB traffic to the operator routers. Those addresses must not overlap with other networks within the site. See Section 2.2.3.2.2 on page 17 for more information.
- **VIP:** VIP networks contain Virtual IP (VIP) addresses that external entities can use in each site to route traffic towards a specific CUDB node. These external entities include CUDB FE applications, provisioning systems, Network

Management Systems (NMSs), and other CUDB nodes. See Section 2.2.3.2.3 on page 18 for more information.

- **SysMGMT**: The SysMGMT network is defined for system administration purposes. See Section 2.2.3.2.4 on page 18 for more information.

As an example, Figure 5 shows a simplified CUDB deployment with three sites. The figure details only a CUDB node and one site, and does not show specific but required non-CUDB equipment. The application FEs and external systems can be connected to public virtual networks as described in the following sections.

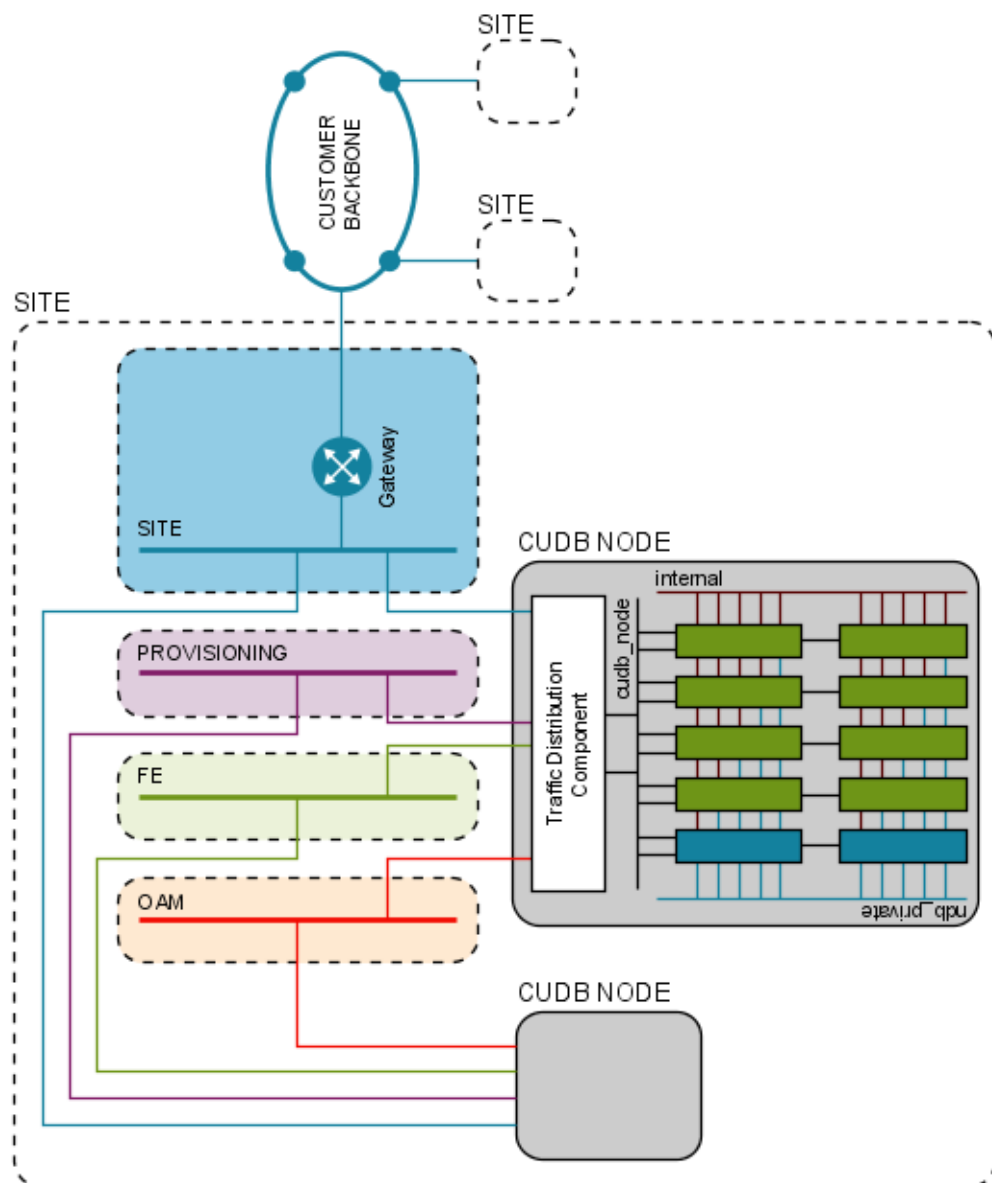


Figure 5 CUDB Deployment Example Showing One CUDB Node Within a Site



Table 1 in Section 2.4 on page 29 provides a summary of the information included in the following sections.

2.2.3.2.1 Internal Networks

Attention!

Only the standard configuration is described, customer specific deployments are not in the scope of this document. Consult the next level of support for more information in case of non-standard configurations.

Internal networks provide internal blade or VM cluster connectivity, and therefore exist in each CUDB node, having exactly the same addresses in each of them. These virtual networks are configured in LDE and (when applicable) in switches and traffic distribution components. All of them are **mandatory**.

If a Provisioning Gateway (PG) is collocated with a CUDB node, other PG-specific networks must be configured.

Due to the different naming requirements, the virtual networks can be identified using different names in LDE configuration, in switches, and in documentation. The following sections explain the virtual networks in detail.

Linux Cluster Networks

This section describes the cluster internal network and, for CUDB systems deployed on native BSP 8100, the boot networks.

LDE Cluster Internal Network

The LDE cluster internal network is dedicated to internal LDE cluster traffic, and is therefore used exclusively by LDE.

In LDE, it is mandatory to name the internal network `internal1`.

The IP range of this virtual network is always on the 192.168.0.0/24 network. It is the same in all CUDB nodes, and is neither visible, nor routable outside the CUDB nodes. Therefore, the following line defining the `internal1` network exists in the `cluster.conf` LDE configuration file:

```
network internal 192.168.0.0/24
```

Boot Networks for CUDB Systems Deployed on Native BSP 8100

Boot networks are specific networks for CUDB nodes based on GEP3/GEP5 blades.



The purpose of these networks is to use 1GbE blade ports instead of the 10GbE ones when booting the second System Controller (SC) and all payload blades on the left and right backplanes.

Note: The 10GbE ports of the GEP5 blades are not DHCP/PXE bootable. Therefore, use the 1GbE ports to load LDE images with the bootL and bootR networks.

The IP range of the bootL virtual network is always on the 192.168.10.0/24 network. In case of bootR, the virtual network IP range is always on the 192.168.11.0/24 network. These IP ranges are the same in all CUDB nodes, and are neither visible nor routable outside the CUDB nodes.

In the `cluster.conf` LDE configuration file, these networks are defined with the following lines:

```
network bootL 192.168.10.0/24
network bootR 192.168.11.0/24
```

application_internal Network

This network is used for CUDB application-specific traffic. The network transfers the following traffic types from the traffic distribution components to the cluster blade or VMs:

- LDAP traffic coming from either external client application FEs, or from other remote CUDB nodes.
- LDAP proxy traffic sent toward other CUDB nodes.
- External OAM orders and requests directed to the SCs.
- Database access from LDAP FEs to CUDB node cluster units (DS Units, or the PLDB).

The IP range of this virtual network is always on the 10.22.0.0/24 network. This is the same in all CUDB nodes, and is neither visible, nor routable outside the CUDB nodes.

In the `cluster.conf` LDE configuration file, this network is defined with the `application_internal` name in its definition:

```
network application_internal 10.22.0.0/24
```

ndb_private Network

The purpose of this network is to transfer the following traffic types:

- Synchronous replication between data nodes belonging to a specific DS Unit or the Processing Layer Database (PLDB) inside a CUDB node.



- Communication between the Network Database (NDB) data nodes and the NDB management nodes of the database cluster. Due to NDB cluster requirements, a specific virtual network for this purpose is required, so the `application_internal` virtual network cannot be used.
- Communication between the `mysqld` process and the NDB data nodes.

The IP range of this virtual network is always on the `10.1.1.0/24` network. This is the same in all CUDB nodes, and is neither visible, nor routable outside the CUDB nodes.

In the `cluster.conf` LDE configuration file, this network is defined with the following line:

```
network ndb_private 10.1.1.0/24
```

2.2.3.2.2 Infrastructure Networks

For CUDB nodes, infrastructure networks provide traffic connectivity, which CUDB supports using the VIP software component.

Traffic separation is managed by the VIP component using separate virtual networks.

- CUDB FE networks for CUDB systems deployed on native BSP 8100:
 - `Cudb_oam_fee`: Used for OAM traffic
 - `Cudb_prov_fee`: Used for CUDB provisioning operations
 - `Cudb_fe_fee`: Used for LDAP traffic proposals
 - `Cudb_site_fee`: Used for replication operations and proxy activities between CUDB nodes within a CUDB system
- CUDB FE networks for CUDB systems deployed on a cloud infrastructure:
 - `vcuodb_oam`: Used for OAM traffic
 - `vcuodb_provisioning`: Used for CUDB provisioning operations
 - `vcuodb_fe`: Used for LDAP traffic proposals
 - `vcuodb_site`: Used for replication operations and proxy activities between CUDB nodes in a CUDB system

Note: To improve the readability of this document, from now on, the relevant infrastructure networks will be referred to as OAM, PROVISIONING, FE, and SITE.

For CUDB systems deployed on native BSP 8100, they refer to `Cudb_oam_fee`, `Cudb_prov_fee`, `Cudb_fe_fee`, and `Cudb_site_fee`, respectively.



For CUDB systems deployed on a cloud infrastructure, they refer to vCUDB_OAM, vCUDB_PROVISIONING, vCUDB_FE, and vCUDB_SITE, respectively.

2.2.3.2.3 VIP Addresses

The VIP addresses available in a CUDB node are as follows:

SITE_VIP(s)

The SITE_VIPs are self-eVIP addresses, used to communicate with CUDB nodes within the same CUDB system through the operator CUDB_SITE network.

OAM_VIP

The OAM_VIP is a self-eVIP address, used to let CUDB nodes communicate with external systems through the operator CUDB_OAM network.

PROVISIONING_VIP

The PROVISIONING_VIP is a self-eVIP address, used to let CUDB nodes communicate with provisioning systems through the operator CUDB_PROVISIONING network.

FE_VIP

The FE_VIP is a self-eVIP address, used to let CUDB nodes communicate with application FE nodes through the operator CUDB_FE network.

2.2.3.2.4 SysMGMT Network

SysMGMT is a mandatory network used for system administration purposes. It provides low-level access with root or administrator privileges to the different components of a CUDB node.

The direct management ports of the CUDB node infrastructure are usually disconnected in normal CUDB operation, and are used only when the CUDB node is installed the first time. In that sense, it is possible to extend the default OAM interface by directly connecting these node infrastructure device management ports to an external network infrastructure, and then build a dedicated system OAM network for better manageability, or for emergency access when the normal OAM access is unavailable for some reason.

The SysMGMT network is connected by using separate links, has regular IP connectivity, and no VIP. The network can be either separated, or included in the customer OAM network during integration. However, take into account that in the latter case, the system addresses become exposed to customer premises. The proper routes to access the SysMGMT network must be applied in the customer network. In any case, it is strongly advised to use the proper security measures.



Note: For CUDB systems deployed on native BSP 8100, the BSP-NBI and SYSMGM virtual networks provide dedicated management access to BSP subsystem and CUDB.

2.2.3.3 External Connectivity for CUDB Systems Deployed on Native BSP 8100

For CUDB systems deployed on native BSP 8100 with GEP3 blades, the supported options for CMX external connectivity are as follows:

- Option 1: 3 uplink-ports configuration. Two types of traffic share the same physical links and use the same port on each CMX. The two types of traffic share the port, but use different virtual networks and tags. This is the default configuration.
- Option 2: 4 uplink-ports configuration. Each type of traffic uses a different physical link and port on each CMX.

In both options, each type of traffic can be configured with any of the available ports. CMXs can be connected either to the Active Patch Panel (APP) of the cabinet to convert from the copper cabling used internally within the cabinet to the optical cabling used externally with the site switches or routers (10GbE ports) or directly to the site switches or routers over copper cabling (1GbE ports).

Figure 6 shows the standard default configuration for 3 uplink-ports, where OAM (om_cn_sp1) and Provisioning (om_cn_sp2) traffic share the same physical links and use the same port on each CMX, and an example of a 4 uplink-ports configuration, where each type of traffic uses a different link and port.

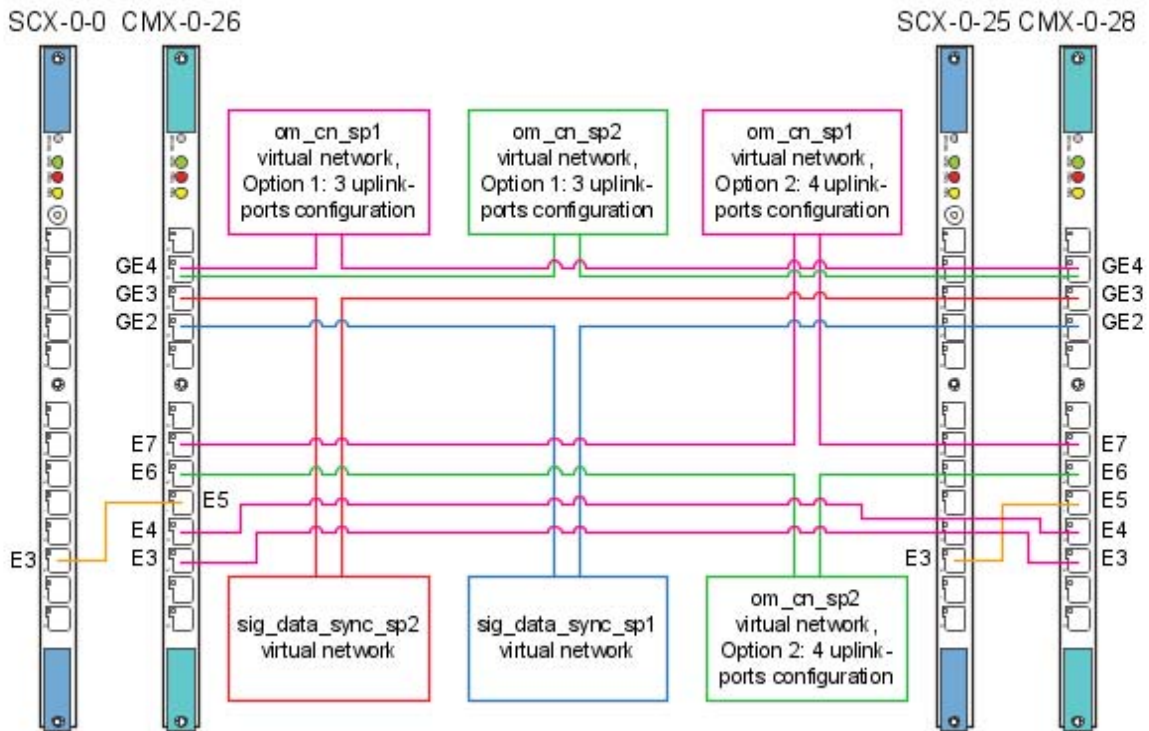


Figure 6 Physical Links Connected to External Site Infrastructure for CUDB Systems Deployed on Native BSP 8100 with GEP3 Blades

For CUDB systems deployed on native BSP 8100 with GEP5 blades, the supported options for CMX external connectivity are as follows:

- Option 1: CMXs are connected to the cabinet's APP to convert from the copper cabling used internally within the cabinet to the optical cabling used externally with the site switches or routers (10GbE ports).
- Option 2: CMXs are directly connected to the site switches or routers over copper cabling (1GbE ports).

As Figure 7 shows, OAM (om_cn_sp1) and Provisioning (om_cn_sp2) share the same physical links and use the same port on each CMX. The two types of traffic share the port, but use different virtual networks and tags. FE (sig_data_sync_sp1) and SITE (sig_data_sync_sp2) traffic share the same physical links and use the same port on each CMX as well.

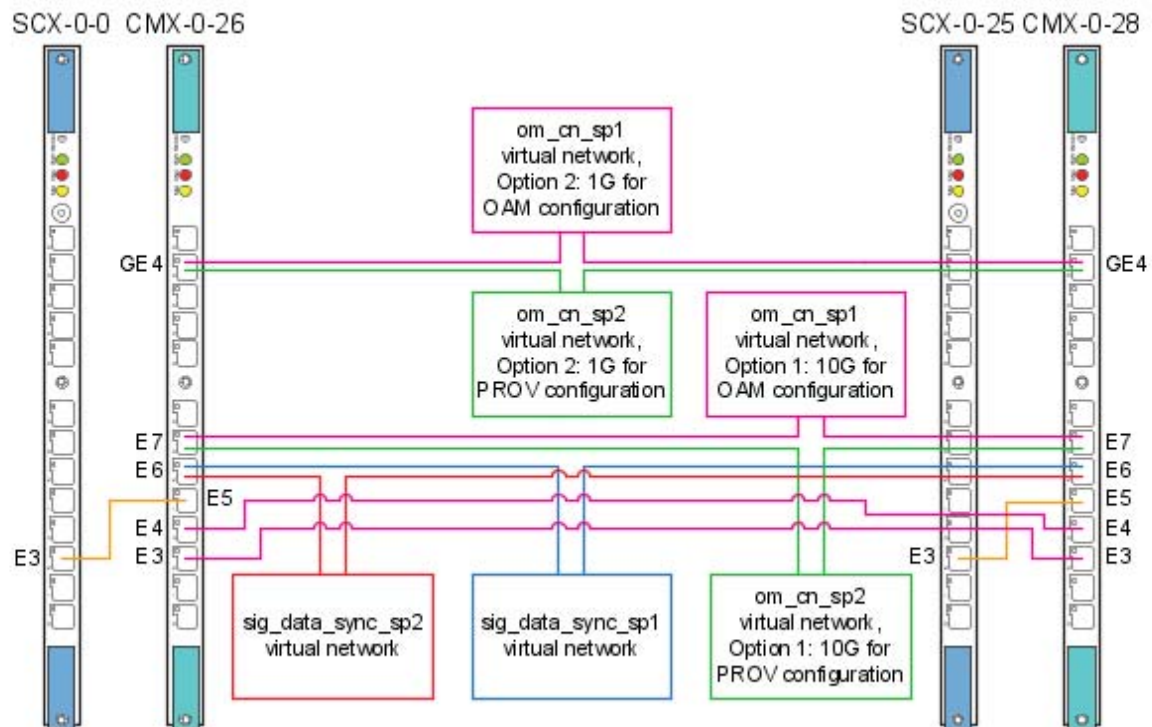


Figure 7 Physical Links Connected to External Site Infrastructure for CUDB Systems Deployed on Native BSP 8100 with GEP5 Blades

Note: Option 1 or Option 2 must be chosen depending on the configuration.

2.2.3.3.1 Uplink Networks for CUDB Systems Deployed on Native BSP 8100

The placement of the CUDB nodes on the customer network must follow traffic separation principles. Traffic separation is used to isolate various traffic types from each other: for example, OAM traffic and user-related traffic must be strictly separated, mostly due to **security** and **overlapping private IP address ranges**, among other reasons. With this approach, the networks described below must exist in the network site where the CUDB node is connected.

Additionally, other networks could also appear related to the specific integration environment of the operator network infrastructure. Such networks are optional, and are not part of the default configuration of a CUDB node. Therefore, they must be deployed as part of an integration project for the customer.

The networks below are defined as virtual networks on the interfaces connected to the external site infrastructure. Assigned IP addresses within these networks are allocated from the operator network IP addressing plan.



None of these virtual networks exist internally in CUDB nodes, so these are not included in the LDE configuration.

CUDB_FE Network

CUDB_FE is a **mandatory** network, used for regular application FE traffic. Application FEs use this network to exchange LDAP application information with CUDB nodes.

CUDB_OAM Network

CUDB_OAM is a **mandatory** network, used to separate OAM traffic from regular application FE traffic. The NMS uses this network to get access to the OAM interfaces of the CUDB nodes.

By default, CUDB provides an OAM interface to manage CUDB node components through a VIP address that redirects OAM traffic to one of the SCs. Once an OAM session is established on the SCs, it is possible to hop on to any blade or VM of the LDE cluster. All OAM tasks are then centralized on the SCs.

CUDB_PROVISIONING Network

CUDB_PROVISIONING is a **mandatory** network, used to separate provisioning traffic from regular application FE traffic. Take into account that provisioning traffic is usually OAM-related traffic, but it must be logically or physically separated. In that sense, a devoted network at site level is included to connect PG nodes and CUDB nodes.

CUDB_SITE Network

CUDB_SITE is a **mandatory** network used to provide connectivity between the CUDB nodes and any other equipment deployed on the site. Additionally, this network contains the site gateway-router that the CUDB nodes use to route traffic to other nodes in other sites. Therefore, all traffic exchanged among CUDB nodes is transferred by this network, and can traverse its gateway if destination is a CUDB node in a remote site.

2.2.3.4 BFD and VRRP for CUDB Systems Deployed on Native BSP 8100

To guarantee the service, a redundancy mechanism is required at CMX level (towards the customer network). When deployed on native BSP 8100, this redundancy mechanism can be configured either at uplink level, that is, independently for each uplink network, or at CUDB node level, that is, with the same redundancy mechanism for all networks. CUDB supports two protocols that can provide redundancy, VRRP and BFD.



VRRP

This protocol is based on a shared IP address and both CMX routers agree which one of them owns the shared IP at a certain time.

The CMX pair is configured in Active-Standby configuration (as preferred configuration) sharing the same IP external addresses by means of VRRP. Nevertheless, VRRP between both routers can be configured to share two IP addresses and get an Active-Active scheme in case the traffic from the site is using ECMP or other mechanism to load balancing traffic over the eVIP gateway routers. In this case, each router will have a primary IP address and also the primary IP of the other router as backup (second IP with less priority in VRRP configuration).

BFD

Contrary to VRRP, the BFD protocol is not based on a shared IP address and the CMX routers do not need to communicate with each other. A different mechanism is used; each router monitors its next hop on the L3 network independently. This next hop usually corresponds to a router in the customer network.

Some configuration is required in customer routers to properly configure BFD.

In the CUDB, there are two possible BFD configuration alternatives:

- Single BFD session
- Multiple BFD session

From the CUDB point of view, there is no significant difference between choosing one protocol or the other. Both VRRP and BFD are able to detect failures in less than a second. Consequently, both of them provide CUDB with a suitable redundancy mechanism. The choice depends on customer requirements.

Note: This document provides the instructions to configure BFD on the CUDB side only. The required actions and configurations in the customer network are not described here.

2.2.3.5

Internal L2 Connectivity for CUDB Systems Deployed on Native BSP 8100

All blades in a subrack are connected to both CMX switches in the subrack through the backplane links.

SCXs are cross link L2 connected through the backplane links.

All blade links use Active-Standby L2 resiliency mechanism based on Linux bonding driver.



The way to connect the CMX routers is to connect both E3 10GbE front ports of the first subrack SCXs to the E5 10GbE front ports of the first subrack CMX routers, as shown in Figure 8.




Figure 8 Connection in a 3-Subrack Configuration



2.2.4 Other Considerations for CUDB Systems Deployed on Native BSP 8100

This section provides some considerations regarding hardware alarms and naming conventions for CUDB systems deployed on native BSP 8100.

2.2.4.1 Hardware Alarms

Hardware alarms are not directly managed by CUDB. Instead, all hardware alarms are forwarded to the NMS by BSP software. Check BSP 8100 documentation to configure NMS as SNMP trap receiver.

Note: This information is provided for reference purposes only.

2.2.4.2 Hardware Naming

Throughout this document, the following naming conventions are used for hardware:

SCX blades:

- SCX-0-0, left SCXB in subrack 0 (SC_2_1, with BSP 8100 software).
- SCX-0-25, right SCXB in subrack 0 (SC_2_2, with BSP 8100 software).
- SCX-1-0, left SCXB in subrack 1 (blade_1_0).
- SCX-1-25, right SCXB in subrack 1 (blade_1_25).
- SCX-2-0, left SCXB in subrack 2 (blade_2_0).
- SCX-2-25, right SCXB in subrack 2 (blade_2_25).

Switches and routers:

- CMX-0-26, left switch/router in subrack 0 (blade_0_26).
- CMX-1-26, left switch/router in subrack 1 (blade_1_26).
- CMX-2-26, left switch/router in subrack 2 (blade_2_26).
- CMX-0-28, right switch/router in subrack 0 (blade_0_28).
- CMX-1-28, right switch/router in subrack 1 (blade_1_28).
- CMX-2-28, right switch/router in subrack 2 (blade_2_28).

2.3 VIP Address Allocation

The following sections describe how to assign the Virtual IP addresses to each CUDB node within each network. These VIPs are used as the communication point from the specified public networks. All of them are network addresses, combined with a protocol and port providing the access point to the CUDB node services provided for customer equipment and other nodes. Some of them are optional, while others correspond to a minimum set of VIPs that each CUDB node requires to be able to provide its basic service. The conditions are explained in the below sections, along with further details.

Table 1 in Section 2.4 on page 29 gives a summary of the information included in the following subsections.



2.3.1 SITE_VIP

SITE_VIP is a **mandatory** address, corresponding to the CUDB node VIP that each node in the CUDB system uses to communicate with the specific node. Therefore, each CUDB node is assigned to one SITE_VIP, and all nodes know their own SITE_VIP as well as the SITE_VIPs assigned to all other nodes.

Any address can be assigned by the operator as long as it is unique across the whole CUDB system, and is routable within the site and among other sites. This routing must be configured in the CUDB_SITE network gateway. Therefore, provided that it is properly setup, it is not required for the SITE_VIP to belong to any network.

Additionally, any outgoing traffic started by the CUDB node towards the SITE_VIP in another node must carry the origin node SITE_VIP as source address. The only exception to this are the multiple SITE_VIPs, where the origin address is any of the SITE_VIPs.

2.3.1.1 Multiple SITE_VIPs

Each CUDB node needs to be configured with the adequate number of SITE_VIP addresses for communication with other CUDB nodes. Each node has a primary SITE_VIP address, and depending on the size of the CUDB system (at least 10 nodes), can also have secondary and source-equivalent SITE_VIP addresses, as well. The formula to calculate the number of IPs needed is as follows:

1. Decrease the number of nodes in the system by 1.
2. Divide the result of the previous step by 8.
3. Finally, round up the result to the next larger integer.

2.3.2 OAM_VIP

OAM_VIP is a **mandatory** address. This address corresponds to the VIP within the CUDB_OAM network that the external CUDB_OAM equipment uses to communicate with the specific CUDB node.

The address must be allocated from the customer network plan, and must belong to the OAM operator network in the site. Any outgoing traffic started by a CUDB node towards the CUDB_OAM network carries the origin node OAM_VIP as the source address. Also, any packet sent from a node towards its OAM_VIP returns to the node.

2.3.3 FE_VIP

FE_VIP is a **mandatory** address. This address corresponds to the VIP within the CUDB_FE network that the external traffic application FEs use to send and receive LDAP traffic to and from the specific CUDB node.



The address must be allocated from the customer network plan and must belong to the CUDB_FE operator network in the site. Any outgoing traffic started by a CUDB node towards the CUDB_FE network carries the origin node FE_VIP as the source address. Also, any packet sent from a node towards its FE_VIP returns to the node.

2.3.4 PROVISIONING_VIP

PROVISIONING_VIP is a **mandatory** address. This address corresponds to the VIP within CUDB_PROVISIONING network that the external provisioning equipment use to send and receive LDAP provisioning traffic to and from the specific CUDB node.

The address must be allocated from the customer network plan, and must belong to the PROVISIONING operator network in the site.

2.3.5 VIP Address Mapping with Configuration Model Attributes

The `CudbLocalNode` and `CudbRemoteNode` configuration model classes (described in [CUDB Node Configuration Data Model Description, Reference \[3\]](#)) include several attributes that are related to the VIP addresses configured in the load balancing solution at every node. These attributes are as follows:

- **Attribute** `cudbVIP` is equal to the `SITE_VIP` address in the `eVIP` component of the `CudbLocalNode` class instance.
- **Attribute** `oamVIP` is equal to the `OAM_VIP` address configured in the `eVIP` component of the `CudbLocalNode` class instance.

In the default deployment described in this document, **attribute `oamVIP` in the `CudbRemoteNode` class instance must be equal to the `SITE_VIP` address configured in the same class**. However, the `oamVIP` attribute can be equal to `OAM_VIP`, if the following conditions are fulfilled:

- The customer requires complete OAM traffic separation, including OAM-like traffic exchanged between CUDB nodes.
 - Connectivity exists among the CUDB_OAM networks in each site, that is, it is possible for any CUDB node to reach OAM interfaces in any other node through the corresponding OAM_VIPs.
- **Attribute** `trafficVIP` is equal to the `FE_VIP` address configured in the `eVIP` component of the `CudbLocalNode` class instance.

In the default deployment described in this document, **attribute `trafficVIP` in the `CudbRemoteNode` class instance must be equal to the `SITE_VIP` address configured in the same class**. However, `trafficVIP` attribute can have a different value, if the following conditions are fulfilled:

- The customer requires traffic separation for the LDAP proxy traffic exchanged between the CUDB nodes.



- A secondary backbone for that purpose is available.

In this case, the `trafficVIP` attribute does not need to correspond to the `FE_VIP` assigned to a particular node, because this backbone would probably be a network different from any `CUDB_FE` network existing in each site. Take into account that this secondary backbone configuration is a customization of the default deployment described in this document: in other words, it requires additional configuration of the `eVIP` component which is not described in this document. Therefore, **this requires a specific customer integration project if requested.**

2.4 Address Plan Summary

Attention!

Only the standard configuration is described, customer specific deployments are not in the scope of this document. Consult the next level of support for more information in case of non-standard configurations.

IPv6 Support

IPv6 configuration is supported for `SITE_VIP(s)`, `OAM_VIP`, `PROVISIONING_VIP`, `FE_VIP`, and `SYSMGMT` addresses. IPv6 is supported only for new installations. Upgrade from IPv4 deployment and dual stack, the combination of IPv4 and IPv6 addresses, are not supported, therefore all addresses must be of the same type. For more information about IPv6 configuration, refer to [CUDB Node Configuration Data Model Description, Reference \[3\]](#).

Table 1 collects the most relevant information for CUDB systems deployed on native BSP 8100.

Table 1 Address Plan Summary for CUDB Systems Deployed on Native BSP 8100

Network Common Name	BSP 8100 Virtual Network Name	CUDB Node Internal Configuration Identifier	Purpose	Type	Net/Mask	Allocated VIPs
lotc	cudb_lde_sp	internal	LDE cluster management.	Private	192.168.0.0/24	N/A
application_internal	cudb_application_internal_sp	application_internal	Provides CUDB node connectivity to handle internal CUDB traffic.	Private	10.22.0.0/24	N/A
ndb_private	cudb_ndb_private_sp	ndb_private	Specific network required for NDB data nodes traffic and management.	Private	10.1.1.0/24	N/A



Table 1 Address Plan Summary for CUDB Systems Deployed on Native BSP 8100

Network Common Name	BSP 8100 Virtual Network Name	CUDB Node Internal Configuration Identifier	Purpose	Type	Net/Mask	Allocated VIPs
Cudb_oam_fee ⁽¹⁾	cudb_om_sp2	N/A	Publishing OAM_VIP	Infrastructure	192.168.218.0/28 ⁽²⁾	One OAM_VIP per CUDB node in the site.
Cudb_prov_fee ⁽¹⁾	cudb_om_sp3	N/A	Publishing PROVISIONING_VIP	Infrastructure	192.168.226.0/28 ⁽²⁾	One PROVISIONING_VIP per CUDB node in the site
Cudb_fe_fee ⁽¹⁾	cudb_sig_sp1	N/A	Publishing FE_VIP	Infrastructure	192.168.216.0/28 ⁽²⁾	One FE_VIP per CUDB node in the site.
Cudb_site_fee ⁽¹⁾	cudb_sig_sp2	N/A	Publishing SITE_VIP	Infrastructure	192.168.217.0/28 ⁽²⁾	One primary SITE_VIP ⁽³⁾ per CUDB node in the site. It is possible to have multiple SITE_VIPs ⁽³⁾ .
CUDB_FE ⁽¹⁾⁽⁴⁾	sig_data_sync_sp1	N/A	Connects the CUDB node with client FE applications.	Uplink	Assigned from the customer network plan.	N/A
CUDB_OAM ⁽¹⁾⁽⁴⁾	om_cn_sp1	N/A	OAM network	Uplink	Assigned from the customer network plan.	N/A
CUDB_PROVISIONING ⁽¹⁾⁽⁴⁾	om_cn_sp2	N/A	Provisioning related traffic.	Uplink	Assigned from the customer network plan.	N/A
CUDB_SITE ⁽¹⁾⁽³⁾	sig_data_sync_sp2	N/A	Connects all CUDB nodes and other solution equipment in the same site, and also includes gateways to other sites.	Uplink	Assigned from the customer network plan.	N/A



Table 1 Address Plan Summary for CUDB Systems Deployed on Native BSP 8100

Network Common Name	BSP 8100 Virtual Network Name	CUDB Node Internal Configuration Identifier	Purpose	Type	Net/Mask	Allocated VIPs
SysMGMT	BSP-NBI and SYSMGMT ⁽¹⁾	sysmgmt	Mandatory network for system administration purposes. Due to BSP limitation, BSP-NBI can be only IPv4 network.	SYSMGMT	Assigned from the customer network plan.	None
NTP Left and NTP Right	cudb_Control-L and cudb_Control-R	ntp_pdl and ntp_pdr	NTP virtual network	Private	192.168.12.0/24 and 192.168.13.0/24	None

(1) This can be IPv4 or IPv6 network.

(2) In case of IPv6 deployment, there is not a fixed address plan for this network. Auto-assigned link-local addresses are used instead.

(3) This network must be visible and routable in the CUDB_SITE virtual network outside the CUDB node. Addresses are selected during site integration.

(4) Networks used to reach the CUDB nodes from the customer networks. Addresses are selected during site integration.

Table 2 collects the most relevant information for CUDB systems deployed on a cloud infrastructure.

Table 2 Address Plan Summary for CUDB Systems Deployed on a Cloud Infrastructure

Network Common Name	CUDB Node Internal Configuration Identifier	Purpose	Type	Net/Mask	Allocated VIPs
lotc	internal	LDE cluster management.	Private	192.168.0.0/24	N/A
application_internal	application_internal	Provides CUDB node connectivity to handle internal CUDB traffic.	Private	10.22.0.0/24	N/A
ndb_private	ndb_private	Specific network required for NDB data nodes traffic and management.	Private	10.1.1.0/24	N/A
vCUDB_FE ⁽¹⁾	N/A	Publishing FE_VIP	Infrastructure	IPv4 example: 192.168.134.0/26 IPv6 example: fd00:1b70:82c8:aaaa::/64	One FE_VIP per CUDB node in the site.
vCUDB_OAM ⁽¹⁾	N/A	Publishing OAM_VIP	Infrastructure	IPv4 example: 192.168.100.0/28 IPv6 example: fd00:1b70:82c8:bbbb::/64	One OAM_VIP per CUDB node in the site.



Table 2 Address Plan Summary for CUDB Systems Deployed on a Cloud Infrastructure

Network Common Name	CUDB Node Internal Configuration Identifier	Purpose	Type	Net/Mask	Allocated VIPs
vCUDB_PROVISIONING ⁽¹⁾	N/A	Publishing PROVISIONING_VIP	Infrastructure	IPv4 example: 192.168.117.0/26 IPv6 example: fd00:1b70:82c8:cccc::/64	One PROVISIONING_VIP per CUDB node in the site
vCUDB_SITE ⁽²⁾	N/A	Publishing SITE_VIP	Infrastructure	IPv4 example: 192.168.151.0/26 IPv6 example: fd00:1b70:82c8:dddd::/64	One primary SITE_VIP ⁽²⁾ per CUDB node in the site. It is possible to have multiple SITE_VIPs ⁽²⁾ .
SysMGMT	sysmgmt	Mandatory network for system administration purposes.	SYSMGMT	Assigned from the customer network plan.	None

(1) Networks used to reach the CUDB nodes from the customer networks. Addresses are selected during site integration.

(2) This network must be visible and routable in the CUDB_SITE virtual network outside the CUDB node. Addresses are selected during site integration.

2.5 Time Synchronization

The time synchronization of the blades or VMs in the CUDB node is essential: it enables the use of proper timestamps in database operations, and helps to correlate log files in case of (fault) tracing activities.

Besides the internal time synchronization of single CUDB node components, all the CUDB nodes in the system must be synchronized as well. Therefore, the configured NTP servers must be the same in all CUDB nodes. In case the availability of the NTP servers is limited to specific CUDB node site locations, all local NTP servers on each CUDB site must be synchronized with each other to facilitate the full CUDB system synchronization.

The configuration of the external NTP servers is performed with the `cluster.conf` LDE configuration file, and is used then by all the blades or VMs in a CUDB node.

Note: For CUDB systems deployed on native BSP 8100, two internal NTP virtual networks are needed (NTP Left and NTP Right). The IP range of these virtual networks is always on the 192.168.12.0/24 and 192.168.13.0/24 networks. This is the same in all CUDB nodes, and are neither visible, nor routable outside the CUDB nodes.

2.6 Traffic Flows

The following sections describe the traffic flows managed by a CUDB node. The sections contain only general information, as specific details can vary depending on the capabilities provided by the routing and switching solutions of the infrastructure used. Therefore, the section does not cover the following topics:



- Quality of service and traffic handling profiles.
- Low level routing details.

2.6.1 Incoming Traffic

This section details traffic flows coming from external entities and other CUDB nodes across a CUDB system. For each traffic flow, the following information is provided:

- **Description** of each traffic flow.
- **Access Point** that the CUDB node exposes as entry point for the traffic flow. This corresponds to the transport address.
- **Enabled on Networks** lists the networks where the traffic flow is enabled, that is, infrastructure networks through which the traffic flow is received and accepted.
- Name of the **Target Pool** in eVIP that handles the traffic inside the CUDB node. The target pool is a group of blades or VMs to which network traffic is distributed.
- **Distribution Policy** in eVIP that states how the traffic is distributed among resources.

2.6.1.1 Incoming OAM Traffic

Table 3 shows different incoming OAM traffic types.

Table 3 Incoming OAM Traffic Types

Traffic Type	Description	Access Point	Enabled on Networks	Target Pool	Distribution Policy
OAM SSH	SSH access to the SCs from the OAM customer network for OAM purposes. The accessed SC corresponds to the one holding the OAM North Bound Interface (NBI) platform. This interface is managed by the Common Operation and Maintenance (COM) CBA component that ensures that the configuration model console (or Command Line Interface, CLI) can be executed in the accessed SC. Also, any other OAM operation not specifically related with the configuration model can be executed in the accessed SC.	TCP port 22 ⁽¹⁾	<ul style="list-style-type: none"> • OAM • SITE • SYSMGMT 	SCs_rr	Round-robin.
OAM NETCONF	Access to NETCONF service for OAM purposes.	TCP port 830 ⁽²⁾	<ul style="list-style-type: none"> • OAM • SITE • SYSMGMT 	SCs_rr	Round-robin.
OAM ESA SNMP	Access to ESA SNMP service for OAM purposes.	UDP port 60 ⁽¹⁾	<ul style="list-style-type: none"> • OAM • SYSMGMT 	SCs_rr	Round-robin.



Table 3 Incoming OAM Traffic Types

Traffic Type	Description	Access Point	Enabled on Networks	Target Pool	Distribution Policy
BSP SSH ⁽²⁾	Access to BSP configuration interface.	TCP port 2024	BSP-NBI	N/A	N/A
CMX SSH ⁽²⁾	Access to CMX (only in case of troubleshooting).	TCP port 22	BSP-NBI	N/A	N/A

(1) The VIP belonging to the CUDB_OAM network allocated from the customer network plan, considered to be the OAM_VIP of the node.

(2) For CUDB systems deployed on native BSP 8100 only.

2.6.1.2 Incoming LDAP Traffic

Table 4 shows different incoming LDAP traffic types.

Table 4 Incoming LDAP Traffic Types

Traffic Type	Description	Access Point	Enabled on Networks	Target Pool	Distribution Policy
CUDB_FE LDAP	Point of access to LDAP FEs to applications from the CUDB_FE customer network.	TCP port 389 ⁽¹⁾	FE	PLs_lc	Least connection
CUDB_FE LDAPS	Point of access to LDAP FEs to applications from the CUDB_FE customer network, using secure LDAP protocol.	TCP port 636 ⁽¹⁾	FE	PLs_lc	Least connection
PG LDAP	Point of access to LDAP FEs from provisioning equipment through the CUDB_PROVISIONING customer network.	TCP port 389 ⁽²⁾	PROVISIONING	PLs_lc	Least connection
PG LDAPS	Point of access to LDAP FEs from provisioning equipment through the CUDB_PROVISIONING customer network, using secure LDAP protocol.	TCP port 636 ⁽²⁾	PROVISIONING	PLs_lc	Least connection
CUDB LDAP	Point of access for LDAP proxy traffic from local and other CUDB nodes.	TCP port 389 ⁽³⁾	SITE	PLs_rr	Round-robin
CUDB LDAP with StartTLS	Point of access for LDAP proxy traffic from local and other CUDB nodes, using LDAP with StartTLS.	TCP port 389 ⁽³⁾	SITE	PLs_rr	Round-robin

(1) The VIP belonging to CUDB_FE network allocated from customer network plan, considered to be the node FE_VIP.

(2) The VIP belonging to CUDB_PROVISIONING network allocated from customer network plan, considered to be the node PROVISIONING_VIP.

(3) If multiple SITE_VIPs exist, the destination is the primary SITE_VIP.

2.6.1.3 Incoming BC Traffic

Table 5 shows different incoming Blackboard Coordinator (BC) traffic types.

Table 5 Incoming BC Traffic Types

Traffic Type	Description	Access Point	Enabled on Networks	Target Pool	Distribution Policy
CUDB BC Access	Handling traffic towards BC server processes running in the node.	TCP port range 9511 :9513 ⁽¹⁾	SITE	BCServers_pool_rr	Round-robin



Table 5 Incoming BC Traffic Types

Traffic Type	Description	Access Point	Enabled on Networks	Target Pool	Distribution Policy
CUDB BC Followers	Handling traffic from the BC server processes acting as followers towards the BC server process acting as leader running in the node.	TCP port range 4511-4513 ⁽¹⁾	SITE	BCServers_pool_rr	Round-robin
CUDB BC Leader Election	Handling traffic for the leader election mechanism of the BC cluster.	TCP port range 5511-5513 ⁽¹⁾	SITE	BCServers_pool_rr	Round-robin

(1) If multiple SITE_VIPs exist, the destination is the primary SITE_VIP.

2.6.1.4 Incoming PLDB Traffic

Table 6 shows different incoming PLDB traffic types.

Table 6 Incoming PLDB Traffic Types

Traffic Type	Description	Access Point	Enabled on Networks	Target Pool	Distribution Policy
CUDB In PLDB Access	Handling traffic towards the database cluster access servers that belong to the PLDB cluster in the node.	TCP port 15000 ⁽¹⁾ This is a recommended value for the CudbP1Group.accessPort attribute (refer to CUDB Node Configuration Data Model Description, Reference [3]). If this rule is not followed, correspondence must be ensured between the pool value and the configuration model.	SITE	PLs_rr_pldb	Round-robin
CUDB In PLDB Repl Master Ch1	Handling traffic towards the database cluster master replication server for the first replication channel that belongs to the PLDB cluster in the node.	TCP port 15001 ⁽¹⁾ This is the recommended value for the CudbP1Group.masterReplicationChannel1Port attribute (refer to CUDB Node Configuration Data Model Description, Reference [3]). If this is not followed, correspondence must be ensured between the pool value and the configuration model.	SITE	PLs_rr_pldb	Round-robin
CUDB In PLDB Repl Master Ch2	Handling traffic towards the database cluster master replication server for the second replication channel that belongs to the PLDB cluster in the node.	TCP port 15002 ⁽¹⁾ This is the recommended value for the CudbP1Group.masterReplicationChannel2Port attribute (refer to CUDB Node Configuration Data Model Description, Reference [3]). If this is not followed, correspondence must be ensured between the pool value and the configuration model.	SITE	PLs_rr_pldb	Round-robin

(1) If multiple SITE_VIPs exist, the destination is the primary SITE_VIP.

2.6.1.5 Incoming DSG Traffic

Table 7 shows different incoming Data Store Unit Group (DSG) traffic types.



Table 7 Incoming DSG Traffic Types

Traffic Type	Description	Access Point	Enabled on Networks	Target Pool	Distribution Policy
CUDB In DSG<g> Access	Handling traffic towards the database cluster access servers belonging to a local data store, provided that the node contains a data store hosting data for DSG number <g>. One or more of this type is defined corresponding to the CudbDsGroup objects defined in the configuration data model (refer to CUDB Node Configuration Data Model Description, Reference [3]).	TCP port1 following the formula: $15000 + (<g> * 10)^{(1)}$ This is the recommended value for the CudbDsGroup.accessPort attribute (refer to CUDB Node Configuration Data Model Description, Reference [3]). If this is not followed, correspondence must be ensured between the pool value and the configuration model.	SITE	PLs_rr_ds<g>	Round-robin
CUDB In DSG<g> Repl Master Ch1	Handling traffic towards the database cluster access server acting as the master for the first replication channel belonging to a local data store, provided that the node contains a data store hosting data for DSG number <g>. One or more of this type is defined corresponding to the CudbDsGroup objects defined in the configuration data model (refer to CUDB Node Configuration Data Model Description, Reference [3]).	TCP port following the formula: $15000 + (<g> * 10) + 1^{(1)}$ This is the recommended value stated in the configuration data model description for the CudbDsGroup.masterReplicationChannel1Port attribute (refer to CUDB Node Configuration Data Model Description, Reference [3]). If this is not followed, correspondence must be ensured between the pool value and the configuration model.	SITE	PLs_rr_ds<g>	Round-robin
CUDB In DSG<g> Repl Master Ch2	Handling traffic towards the database cluster access server acting as the master of the second replication channel belonging to a local data store, provided that the node contains a data store hosting data for DSG number <g>. One or more of this type is defined corresponding to the CudbDsGroup objects defined in the configuration data model (refer to CUDB Node Configuration Data Model Description, Reference [3]).	TCP port following the formula: $15000 + (<g> * 10) + 2^{(1)}$ This is the recommended value stated in the configuration data model description for the CudbDsGroup.masterReplicationChannel2Port attribute (refer to CUDB Node Configuration Data Model Description, Reference [3]). If this is not followed, correspondence must be ensured between the pool value and the configuration model.	SITE	PLs_rr_ds<g>	Round-robin

(1) If multiple SITE_VIPs exist, the destination is the primary SITE_VIP.

2.6.2 Outgoing Traffic

This section describes the details of the traffic originated in a CUDB node towards external entities and other CUDB nodes. Each subsection depends on the type of traffic originated from CUDB and on the outgoing gateway used.

For each traffic flow, the following information is provided:

- **Description** of each traffic flow.
- **Destination IP Address** that is the IP address configured as a destination.
- **Destination Port** that is the network or equipment receiving the traffic and, when applicable, the UDP or TCP port.
- **Source Address** that is the required source address to set in outgoing packets.
- **Gateway** to use in case the traffic must traverse its immediate receiving network.



2.6.2.1 Outgoing OAM Traffic

Table 8 shows different outgoing OAM traffic types.

Table 8 Outgoing OAM Traffic Types

Traffic Type	Description	Destination IP Address	Destination Port	Source Address	Gateway
OAM Out SNMP	SNMP traffic generated from the node towards the NMS acting as trap collector.	IP address stated in the ESA V3 trap forwarding table configuration file (/home/cudb/oam/faultMgmt/config/masterAgent/trapDestCfg.xml). ⁽¹⁾	Port stated in ESA V3 trap forwarding table configuration file (/home/cudb/oam/faultMgmt/config/masterAgent/trapDestCfg.xml). ⁽¹⁾	OAM_VIP	OAM network
OAM Out NTP	NTP requests	NTP servers configured in the /cluster/etc/cluster.conf file, and with the ntp parameter.	UDP port 123	OAM_VIP	OAM network
OAM	Authentication LDAP queries sent to an external LDAP server. ⁽²⁾	IP addresses stated in the primaryServer and secondaryServer attributes in the corresponding CudbExternalAuthServer model class instance.	TCP ports 389 or 636	OAM_VIP	OAM network

(1) It is possible to configure more than one NMS acting as trap collector.

(2) Outgoing traffic type of the CUDB OAM Centralized Authentication System Support function.

2.6.2.2 Outgoing PG Traffic

Table 9 shows outgoing PG traffic types.

Table 9 Outgoing PG Traffic Types

Traffic Type	Description	Destination IP Address	Destination Port	Source Address	Gateway
CUDB Out PG <port>	Outgoing provisioning JMX notifications towards the provisioning gateways. There is an outgoing traffic flow of this type per port in the CudbProvisioningGatewayConfig configuration model class instance, in the pgNodeIpAddresses multi-valued attribute. For more information, refer to CUDB Node Configuration Data Model Description, Reference [3].	Each IP address configured in the configuration model.	Each port configured in the configuration model.	OAM_VIP	OAM network
CUDB Out Provisioning Assurance <port>	Outgoing HTTP traffic towards the provisioning gateways. There is an outgoing traffic flow of this type per port in the CudbProvGatewayEndPoint configuration model class instance, in the replayRequestURL and replayStatusURL attributes. For more information, refer to CUDB Node Configuration Data Model Description, Reference [3].	Each IP address configured in the configuration model.	Each port configured in the configuration model.	OAM_VIP	OAM network



2.6.2.3 Outgoing FE Traffic

Table 10 shows outgoing FE traffic.

Table 10 Outgoing FE Traffic Type

Traffic Type	Description	Destination IP Address	Destination Port	Source Address	Gateway
CUDB Out NOTIFICATIONS <port>	Outgoing traffic notifications towards application FEs. There is an outgoing flow of this type per port in the CudbNotificationEndPoint configuration model class instance.	IP address or host name stated in the URI attribute in the corresponding CudbNotificationEndPoint class instance.	Port in the URI attribute in the corresponding CudbNotificationEndPoint class instance.	FE_VIP	FE network

2.6.2.4 Outgoing LDAP Traffic

Table 11 shows different outgoing LDAP traffic types.

Table 11 Outgoing LDAP Traffic Types

Traffic Type	Description	Destination IP Address	Destination Port	Source Address	Gateway
CUDB Out LDAP	LDAP proxy traffic towards other CUDB nodes.	Any IP (0.0.0.0). The specific destination is a remote node SITE_VIP.	TCP port 389	SITE_VIP ⁽¹⁾	SITE network
CUDB Out LDAP with StartTLS	LDAP with StartTLS proxy traffic towards other CUDB nodes.	Any IP (0.0.0.0). The specific destination is a remote node SITE_VIP.	TCP port 389	SITE_VIP ⁽¹⁾	SITE network

(1) If multiple SITE_VIPs exist, the source can be any configured VIP.

2.6.2.5 Outgoing BC Traffic

Table 12 shows different outgoing BC traffic types.

Table 12 Outgoing BC Traffic Types

Traffic Type	Description	Destination IP Address	Destination Port	Source Address	Gateway
CUDB Out BC Access	Outgoing traffic towards the BC server processes running in other CUDB nodes.	Any IP (0.0.0.0). The specific destination is a remote node SITE_VIP.	TCP port range 9511:9513	SITE_VIP ⁽¹⁾	SITE network
CUDB Out BC Followers	Outgoing traffic from the BC servers processes acting as followers towards the BC server process acting as leader running in other CUDB node.	Any IP (0.0.0.0). The specific destination is a remote node SITE_VIP.	TCP port range 4511:4513	SITE_VIP ⁽¹⁾	SITE network
CUDB Out BC Leader Election	Outgoing traffic for the leader election mechanism of the BC cluster.	Any IP (0.0.0.0). The specific destination is a remote node SITE_VIP.	TCP port range 5511:5513	SITE_VIP ⁽¹⁾	SITE network

(1) If multiple SITE_VIPs exist, the source is the primary SITE_VIP.



2.6.2.6 Outgoing SSH Traffic

Table 13 shows outgoing SSH traffic.

Table 13 Outgoing SSH Traffic Type

Traffic Type	Description	Destination IP Address	Destination Port	Source Address	Gateway
CUDB Out SSH	Outgoing SSH connections towards other CUDB nodes.	Any IP (0.0.0.0). The specific destination is a remote node SITE_VIP.	TCP port 22	SITE_VIP ⁽¹⁾	SITE network

(1) If multiple SITE_VIPs exist, the source is the primary SITE_VIP.

2.6.2.7 Outgoing PLDB Traffic

Table 14 shows different outgoing PLDB traffic types.

Table 14 Outgoing PLDB Traffic Types

Traffic Type	Description	Destination IP Address	Destination Port	Source Address	Gateway
CUDB Out PLDB Access	Outgoing connections towards the database cluster access servers belonging to the PLDB in other CUDB nodes.	Any IP (0.0.0.0). The specific destination is a remote node SITE_VIP.	TCP port 15000	SITE_VIP ⁽¹⁾	SITE network
CUDB Out PLDB Repl Master Ch1	Outgoing connections towards the database cluster master server for the first PLDB replication channel in other CUDB nodes.	Any IP (0.0.0.0). The specific destination is a remote node SITE_VIP.	TCP port 15001	SITE_VIP ⁽¹⁾	SITE network
CUDB Out PLDB Repl Master Ch2	Outgoing connections towards the database cluster master server for the second PLDB replication channel in other CUDB nodes.	Any IP (0.0.0.0). The specific destination is a remote node SITE_VIP.	TCP port 15002	SITE_VIP ⁽¹⁾	SITE network

(1) If multiple SITE_VIPs exist, the source is the primary SITE_VIP.

2.6.2.8 Outgoing DSG Traffic

Table 15 shows different outgoing DSG traffic types.

Table 15 Outgoing DSG Traffic Types

Traffic Type	Description	Destination IP Address	Destination Port	Source Address	Gateway
CUDB Out DSG<g> Access	Outgoing connections towards the database cluster access servers belonging to data stores in other CUDB nodes hosting data for DSG <g>.	Any IP (0.0.0.0). The specific destination is a remote node SITE_VIP.	TCP port following the formula: 15000 + (<g> * 10)	SITE_VIP ⁽¹⁾	SITE network



Table 15 Outgoing DSG Traffic Types

Traffic Type	Description	Destination IP Address	Destination Port	Source Address	Gateway
CUDB Out DSG<g> Repl Master Ch1	Outgoing connections towards the database cluster server in other CUDB nodes acting as the first replication channel for DSG <g>.	Any IP (0.0.0.0). The specific destination is a remote node SITE_VIP.	TCP port following the formula: $15000 + (<g> * 10) + 1$	SITE_VIP ⁽¹⁾	SITE network
CUDB Out DSG<g> Repl Master Ch2	Outgoing connections towards the database cluster server in other CUDB nodes acting as the second replication channel for DSG <g>.	Any IP (0.0.0.0). The specific destination is a remote node SITE_VIP.	TCP port following the formula: $15000 + (<g> * 10) + 2$	SITE_VIP ⁽¹⁾	SITE network

(1) If multiple SITE_VIPs exist, the source is the primary SITE_VIP.

2.6.2.9 CUDB OAM Centralized Security Event Logging Traffic

Table 16 shows the outgoing traffic type of the CUDB OAM Centralized Security Event Logging function.

Table 16 OAM Centralized Security Event Logging Outgoing Traffic Type

Traffic Type	Description	Destination IP Address	Destination Port	Source Address	Gateway
OAM	Security logging events sent to an external server.	IP address stated in the externalLogServerIp attribute in the corresponding CudbExternalLogServer model class instance.	TCP port stated in the externalLogServerPort attribute in the corresponding CudbExternalLogServer model class instance.	CUDB OAM	OAM Network

2.7 Routing

Routing is configured in the gateways belonging to each network existing in the site. It is out of the scope of this document to define the low-level details of this configuration, but the following constraints must be taken into account:

- Each network in the described solution contains a gateway configured to route traffic addressed to a VIP towards the corresponding node.
- With the exception of the CUDB_SITE gateway, no other gateways are needed to provide routing to other sites, according to the solution described in this document.
- The CUDB_SITE gateway is able to route incoming traffic from other CUDB nodes in other sites to local CUDB nodes, based on their corresponding SITE_VIPs.



- The system has four VIP addresses divided into different virtual networks, making up the traffic separation both in the system and outwards. Each virtual network has one unique interface in the nodes. The traffic separation must be met by configuring routing accordingly. This is done by combining a default route and one or several static routes in each node. The eVIP framework configures these routes.

Note: If multiple SITE_VIP addresses are configured in the node, the system has more than four VIP addresses.

2.8 Firewall Configuration

When an external firewall is used, it must allow both incoming and outgoing traffic flows to traverse any firewall (as described in Section 2.6.1 on page 33 and Section 2.6.2 on page 36, respectively).

Besides the above traffic, the proprietary protocols used in CUDB for supervision and replication flows carried over the CUDB system, and the external VLAN/VPN must also be allowed in the external firewalls protecting the CUDB external domain.

Note: SOAP-based CUDB notifications traffic must also be allowed in the external firewalls. This traffic does not use a fixed port, but the ones configured in the CUDB node. Refer to [CUDB Node Configuration Data Model Description, Reference \[3\]](#) for further information.

Depending on the application, more protocols may be required to pass the firewall.

Table 17 shows the traffic and ports that must be allowed in the external firewalls.

Table 17 Traffic Flows Allowed in External Firewall

Traffic	Port	Networks
LDAP/LDAPS	TCP/389 (In/Out) TCP/636 (In/Out)	CUDB_FE CUDB_SITE ⁽¹⁾ CUDB_OAM CUDB_PROVISIONING
NTP	UDP/123 (Out)	CUDB_OAM
Database Cluster	TCP/15000-17552 (In/Out) ⁽²⁾	CUDB_SITE
SSH	TCP/22 (In/Out)	CUDB_SITE CUDB_OAM
NETCONF	TCP/830 (In)	CUDB_SITE CUDB_OAM
BC Servers Access	TCP/9511-9513 (In/Out)	CUDB_SITE
BC Servers Followers	TCP/4511-4513 (In/Out)	CUDB_SITE
BC Servers Leader Election	TCP/5511/5513 (In/Out)	CUDB_SITE
SNMP	UDP/<nms_port>(Out) ⁽³⁾	CUDB_OAM



Table 17 Traffic Flows Allowed in External Firewall

Traffic	Port	Networks
Notifications to PG (JMX protocol)	TCP/<pg_ports>(Out) ⁽⁴⁾	CUDB_OAM ⁽⁵⁾
Provisioning Assurance towards PG	TCP/<pg_ports>(Out) ⁽⁴⁾	CUDB_OAM ⁽⁵⁾
Notifications to applications (SOAP protocol)	TCP/<defined EP ports>(Out) ⁽⁶⁾	CUDB_FE
IGMP protocols		All networks

(1) If multiple SITE_VIPs exist, traffic must be allowed for LDAP ports in all CUDB_SITE virtual networks if there is more than one.

(2) Considering the highest supported DS group number as 255, taking 15000 as base port, and taking into account the highest port corresponding to the second replication channel, the formula can be $15000 + (255 * 10) + 2$.

(3) NMS port configured for forwarding SNMP traps. See outgoing traffic description in Table 8.

(4) One port per configured PG. See outgoing traffic description in Table 9.

(5) Through backbone when PG is in different site locations than any of the CUDB nodes.

(6) One port per configured notification end point. See outgoing traffic description in Table 10.



3 Appendix: Quality of Service

By default, no Quality of Service (QoS) criteria is applied in CUDB to any network traffic flow, but it is provided optionally on request. In order to use QoS, traffic is marked with the proper Differentiated Services Code Point (DSCP). Based on this marking, network elements can prioritize traffic.

For the recommended DSCP marking values, see Table 18.

Table 18 Recommended DSCP Marking Values

Network	Traffic Flow	Port	DSCP Value
SYSMGMT OAM_VIP	CUDB_OAM Out NTP	UDP port 123	48
SYSMGMT OAM_VIP SITE_VIP	CUDB_OAM In/Out SSH	TCP port 22	16
OAM_VIP	CUDB_OAM In SNMP CUDB_OAM Out SNMP	UDP port 60 UDP port 162	32
OAM_VIP SITE_VIP	CUDB_OAM In NETCONF	TCP port 830	16
OAM_VIP	CUDB_OAM Out PG Notifications	TCP ports 8994/8099	40
FE_VIP	CUDB_FE In LDAP CUDB_FE In LDAPS	TCP ports 389 TCP port 636	40
FE_VIP	CUDB_FE Out SOAP	TCP port <defined EP ports>	40
SITE_VIP	CUDB_FE In/Out Database Cluster	TCP port range 15000:17552	8
SITE_VIP	CUDB_SITE In/Out LDAP CUDB_SITE In/Out LDAPS	TCP port 389 TCP port 636	40
SITE_VIP	CUDB_SITE In/Out BC Access CUDB_SITE In/Out Followers CUDB_SITE In/Out LeaderElection	TCP port range 9511:9513 TCP port range 4511:4513 TCP port range 5511:5513	48
PROVISIONING_VIP	CUDB_PROVISIONING In LDAP CUDB_PROVISIONING In LDAPS	TCP port 389 TCP port 636	40

Note: In CUDB, DSCP code is not applied to any traffic not included in Table 18.





Glossary

For the terms, definitions, acronyms, and abbreviations used in this document, refer to [CUDB Glossary of Terms and Acronyms, Reference \[4\]](#).





Reference List

CUDB Documents

- [1] CUDB Technical Product Description
- [2] CUDB Node Hardware Description
- [3] CUDB Node Configuration Data Model Description
- [4] CUDB Glossary of Terms and Acronyms

Other Ericsson Documents

- [5] eVIP Management Guide