

Virtual Library Extension

Planning Guide

Version 1.3



Revision 04
Part Number: E39444-04

Submit comments about this document to STP_FEEDBACK_US@ORACLE.COM.

Planning Guide

E39444-04

Copyright © 2013 Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related software documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle USA, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications which may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure the safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle is a registered trademark of Oracle Corporation and/or its affiliates. Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. UNIX is a registered trademark licensed through X/Open Company, Ltd.

This software or hardware and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

Table of Contents

List of Tables	5
Preface	7
Audience	7
Documentation Accessibility	7
Access to Oracle Support	7
Supplementary VLE Documents	7
1 Physical Site Planning	9
Site Evaluation – External Considerations	9
Site Evaluation – Internal Considerations	11
VLE Environmental Specifications	12
Requirements for Transferring the VLE Point to Point	14
Requirements for Installing the VLE	15
Data Center Safety	17
Site Power Distribution Systems	18
HVAC Requirements	22
Environmental Requirements and Hazards	22
2 VLE Planning	23
Satisfying Mainframe Host Software Requirements	23
Satisfying Network Infrastructure Requirements	23
Satisfying Oracle Switch Hardware Requirements	25
Satisfying Serviceability Requirements	26
ASR Configuration	27
Determining VLE Configuration Values.....	28
Determining Values for the Configuration Scripts	28
Determining Values for setup_network	29
Determining Values for setup_vle_node	30
Determining Values for Port Card Configuration	31
Determining VMVC Range Configuration Values	36
Planning for Encryption	38
Planning for Deduplication	39
Planning for Link Aggregation.....	42
Benefits of Link Aggregation	42

Link Aggregation Requirements	43
Switch Configuration	43
10GigE Port Aggregations	44
Monitoring Aggregations	44
Types of VLE Aggregations	45
A VLE Configuration Examples	47
VLE GUI Configuration Screens	48
Port Card Configuration Tab	48
Create New Route/Destination dialog boxes	49
Example 1: One VTSS Connected to One VLE	50
Example 2: Four VTSSs Connected to One VLE	52
Example 3: VLE to VLE Copy	55
Example 4: One VTSS Connected to a 3-Node VLE	58
B VLE Link Aggregation Examples	61
VLE GUI Configuration Screens	62
Port Aggregations tab	62
Example 1: Single-Node VLE Connected to Two VSM5s	63
Example 2: Two-Node VLE Connected to Two VTSSs	64
C Controlling Contaminants	65
Environmental Contaminants	65
Required Air Quality Levels	65
Contaminant Properties and Sources	66
Operator Activity	67
Hardware Movement	67
Outside Air	67
Stored Items	67
Outside Influences	67
Cleaning Activity	68
Contaminant Effects	68
Physical Interference	68
Corrosive Failure	68
Shorts	69
Thermal Failure	69
Room Conditions	69
Exposure Points	70
Filtration	71
Positive Pressurization and Ventilation	72
Cleaning Procedures and Equipment	72
Daily Tasks	73
Weekly Tasks	73
Quarterly Tasks	74
Biennial Tasks	74
Activity and Processes	75

List of Tables

TABLE 1-1	VLE Floor Loading Specifications	16
TABLE 1-2	Raised Flooring Horizontal Force Chart.....	16
TABLE 1-3	Source Power Requirements for VLE Equipment.....	20
TABLE 2-4	ELS Supporting PTFs for VLE	23
TABLE 2-5	CAM Configuration Information	26
TABLE 2-6	Notification Setup - Email Configuration Options / ConfCollectStatus	27
TABLE 2-7	VLE Effective Capacities - Maximum VMVCs Per Node	37
TABLE A-8	Example 1 Configuration Values	51
TABLE A-9	Example 2 Configuration Values	53
TABLE A-10	Example 3 Configuration Values	56
TABLE A-11	Example 4 Configuration Values	59

Preface

Audience

This publication is intended for Oracle or customer personnel responsible for doing site planning for Oracle's StorageTek Virtual Library Extension (VLE).

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

Access to Oracle Support

Oracle customers have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

Supplementary VLE Documents

You are entitled to the following supplementary documents:

- *VLE SLA and Standard Firmware Entitlement*
- *VLE Documentation Third Party Licenses and Notices*
- *Written Offer to Provide Certain Source Code*
- *VLE Safety Guide*

Physical Site Planning

This chapter provides information about activities designed to ensure the site is equipped to accommodate the power, safety, environmental, HVAC, and data handling requirements of VLE system equipment.

Key site readiness planning considerations include, but are not limited to:

- Site surveys to evaluate and eliminate or mitigate factors which could negatively affect delivery, installation, and operation of VLE system equipment.
- A plan for the layout and location of VLE system equipment and cabling that allows for efficient use and easy maintenance, plus adequate space and facilities for Oracle support personnel and their equipment.
- Facilities construction that provides an optimum operating environment for VLE system equipment and personnel, as well as safe flooring and protection from fire, flooding, contamination, and other potential hazards.
- Scheduling of key events and task completion dates for facilities upgrades, personnel training, and delivery, implementation, installation, testing, and certification activities.

Customers ultimately are responsible for ensuring that their site is physically prepared to receive and operate VLE system equipment, and that the site meets the minimum specifications for equipment operation as detailed in this guide.

Site Evaluation – External Considerations

Before delivery of VLE system equipment, a readiness planning team should identify and evaluate all external site factors that present existing or potential hazards, or which could adversely affect delivery, installation, or operation of the system.

External factors that should be evaluated include:

- Reliability and quality of electrical power provided by the local utility, backup power generators, and uninterruptible power supplies (UPSs), and so on,
- Proximity of high-frequency electromagnetic radiation sources (for example, high-voltage power lines; television, radio, and radar transmitters)
- Proximity of natural or man-made floodplains and the resultant potential for flooding in the data center

- Potential effects of pollutants from nearby sources (for example, industrial plants). For more information, see [“Controlling Contaminants” on page 65](#).

If any existing or potential negative factors are discovered, the site readiness planning team should take appropriate steps to eliminate or mitigate those factors before VLE system equipment is delivered. Oracle Global Services offers consultation services and other assistance to identify and resolve such issues. Contact your Oracle account representative for more information.

Site Evaluation – Internal Considerations

Before delivery of VLE system equipment, a readiness planning team should identify and evaluate all internal site factors that present existing or potential hazards, or which could adversely affect delivery, installation, or operation of the system. Internal factors that should be evaluated include:

- Structural dimensions, elevator capacities, floor-load ratings, ramp inclines, and other considerations when transferring equipment point-to-point between the delivery dock, staging area, and data center installation site, as described in:
 - [“VLE Environmental Specifications” on page 12](#)
 - [“Requirements for Transferring the VLE Point to Point” on page 14](#)
- [“Requirements for Installing the VLE” on page 15](#), which describes floor construction and loading requirements.
- Data center safety system design features and capabilities as described in [“Data Center Safety” on page 17](#).
- Site power system(s) design and capacity as described in [“Site Power Distribution Systems” on page 18](#)
- Data center HVAC design features and capabilities as described in [“HVAC Requirements” on page 22](#)
- Environmental requirements as described in:
 - [“Environmental Requirements and Hazards” on page 22](#)
 - [“Controlling Contaminants” on page 65](#)

If any existing or potential negative factors are discovered, the site readiness planning team should take appropriate steps to eliminate or mitigate those factors before VLE system equipment is delivered. Oracle Global Services offers consultation services and other assistance to identify and resolve such issues. Contact your Oracle account representative for more information.

VLE Environmental Specifications

Note – Statistics for power and cooling data are approximate due to variations in data rates and the number of operations occurring.

Base Configuration

The base configuration consists of a Sun Server X2-4, with two 500GB Internal SATA drives, four quad port Gigabit Ethernet NICs, two dual port 10Gb Ethernet cards, two Thebe dual port SAS HBAs, two J4410 JBODs, each populated with 24 3TB SAS HDD, and the SunRack II 1242 Cabinet with dual 10KVA PDUs. The only option is additional capacity, in increments of two JBODs, up to a maximum total of eight.

Capacity

- Base Capacity - Native 82 TB, Effective 330 TB
- Max Capacity - Native 330 TB, Effective 1.32 PB

VLE Overall Dimensions - SunRack II 1242 Cabinet (inches)

- Height - 78.7
- Width - 23.6
- Depth - 47.2

Service Clearance (inches)

- Top - 36

Note – 36 inches is the generic Sun Rack II specification. VLE does not require access through the top except for power cables.

- Front - 36
- Rear - 36

Weight (Pounds, Fully Populated with 8 JBODs)

Breakdown:

- Server - 85
- Cabinet - 332
- Each JBOD - 110.25
- 8 JBODs - 882

Note – Each JBOD - 110.25

- Total Weight - 1299
- Total Weight plus shipping material - 1570

Power ((Base Watts – 2232, Max Watts – 5436)

Breakdown:

- Server - 1164 (peak) 722 (Idle)
- Each JBOD- 534 (peak) 389 (Idle)
- 8 JBODs- 4272 (peak) 3112 (Idle)
- Total Power- 5436 (peak) 3834 (Idle)

HVAC ((Base BTU/HR – 7618, Max BTU/Hr – 18556)

Breakdown:

- Server - 3972 (peak) 2464 (Idle)
- Each JBOD - 1823 (peak) 1328 (Idle)
- 8 JBODs - 14584 (peak) 10624 (Idle)
- Maximum Heat - 18556 (peak) 13088 (Idle)

Requirements for Transferring the VLE Point to Point

Site conditions must be verified to ensure all VLE system equipment can be safely transported between the delivery dock, staging area, and data center without encountering dimensional restrictions, obstructions, or safety hazards, or exceeding rated capacities of lifting and loading equipment, flooring, or other infrastructure. Conditions that must be verified are described below.

Structural Dimensions and Obstructions

Dimensions of elevators, doors, hallways, and so on, must be sufficient to allow unimpeded transit of VLE cabinets (in shipping containers, where appropriate) from the delivery dock to the data center installation location. See [“VLE Overall Dimensions - SunRack II 1242 Cabinet \(inches\)” on page 12](#) for VLE cabinet-dimension details.

Elevator Lifting Capacities

Any elevators that will be used to transfer VLE cabinets must have a certified load rating of at least 1000 kg (2200 lbs.). This provides adequate capacity to lift the heaviest packaged, fully-populated VLE cabinet, a pallet jack (allow 100 kg/220 lbs.), and two persons (allow 200 kg/440 lbs.). See [“Weight \(Pounds, Fully Populated with 8 JBODs\)” on page 12](#) for additional cabinet-weight details.

Ramp Inclines

To prevent VLE cabinets from tipping on ramps while being moved from point to point, the site engineer or facilities manager must verify the incline angle of all ramps in the transfer path. Inclines cannot exceed 10 degrees (176 mm/m; 2.12 in./ft.).

Requirements for Installing the VLE

The following sections describe requirements for installing the VLE.

Floor Construction Requirements

VLE system equipment is designed for use on either raised or solid floors. Carpeted surfaces are not recommended since these retain dust and contribute to the buildup of potentially damaging electrostatic charges. A raised floor is preferable to a solid floor since it permits power and data cables to be located safely away from floor traffic and other potential floor-level hazards.

Floor-Load Ratings

Solid floors, raised floors, and ramps located along the transfer path for VLE cabinets must be able to withstand concentrated and rolling loads generated by the weight of a populated cabinet, equipment used to lift a cabinet (for example, a pallet jack), and personnel who are moving the cabinet from point to point.

Raised floor panels located along a transfer path must be able to resist a concentrated load of 620 kg (1365 lbs.) and a rolling load of 181 kg (400 lbs.) anywhere on the panel, with a maximum deflection of 2 mm (0.08 in.). Raised floor pedestals must be able to resist an axial load of 2268 kg (5000 lbs.). See [“Floor Loading Requirements” on page 15](#) for additional floor-loading details.

When being moved from one location to another, a VLE cabinet generates roughly twice the floor load as in a static state. Using 19 mm (0.75 in.) plywood along a transfer path reduces the rolling load produced by a cabinet.

Floor Loading Requirements

Warning – Exceeding recommended raised-floor loads can cause a floor collapse, which could result in severe injury or death, equipment damage, and infrastructure damage. It is advisable to have a structural engineer perform a floor-load analysis before beginning installation of VLE system equipment.

Caution – When being moved, a VLE cabinet creates almost twice the floor load as when static. To reduce floor load and stress, and the potential for damage or injury when moving a VLE (for example, during installation), consider using 19 mm/0.75 in. plywood on the floor along the path where the cabinet will be moved.

Flooring with an overall (superimposed) load rating of 490 kg/m² (100 lbs./ft²) is recommended. If floors do not meet this rating, a site engineer or facilities manager must consult the floor manufacturer or a structural engineer to calculate actual loads and determine if the weight of a particular VLE system configuration can be safely supported.

Specific information on floor construction requirements is available from the VLE Backline Support group.

Floor Loading Specifications and References

TABLE 1-1 VLE Floor Loading Specifications

Basic Floor Load*	Maximum Superimposed Floor Load #
695 kg/m ² (142 lbs./ft ²)	462 kg/m ² (94lbs./ft ²)

Note –

- * Load over footprint surface area (7093.7 cm²/1099.5 in²) of an unpackaged VLE cabinet, with a maximum weight of 590 kg/1299 lbs., that is, a VLE with 192 array disk drives.
- # Assumes minimum Z+Z axis dimension of 185.3 cm/73.0 in. (that is, cabinet depth 77.1 cm/30.4 in. + front service clearance of 54.1 cm/21.3 in. + rear service clearance of 54.1 cm/21.3 in.), minimum X+X axis dimension of 104.9 cm/41.2 in. (that is, cabinet width 92.1 cm/36.3 in. + left clearance of 6.4 cm/2.5 in. + right clearance of 6.4 cm/2.5 in.).

Raised-Floor Lateral Stability Ratings

In areas of high earthquake activity, the lateral stability of raised floors must be considered. Raised floors where VLE system equipment is installed must be able to resist the horizontal-stress levels shown in [TABLE 1-2 on page 16](#).

TABLE 1-2 Raised Flooring Horizontal Force Chart

Seismic Risk Zone	Horizontal Force (V) Applied at Top of Pedestal
1	13.5 kg / 29.7 lbs
2A	20.2 kg / 44.6 lbs
2B	26.9 kg / 59.4 lbs
3	40.4 kg / 89.1 lbs
4	53.9 kg / 118.8 lbs

Note – Horizontal forces are based on the 1991 Uniform Building Code (UBC) Sections 2336 and 2337, and assume minimum operating clearances for multiple VLE cabinets. Installations in areas not covered by the UBC should be engineered to meet seismic code provisions of the local jurisdiction.

Raised-Floor Panel Ratings

Raised floor panels must be able to resist a concentrated load of 590 kg (1299 lbs.) and a rolling load of 181 kg (400 lbs.) anywhere on the panel with a maximum deflection of 2 mm (0.08 in.). Perforated floor panels are not required for VLE system equipment, but if used must comply with the same ratings.

Raised-Floor Pedestal Ratings

Raised floor pedestals must be able to resist an axial load of 2268 kg (5000 lbs.). Where floor panels are cut to provide service access, additional pedestals may be required to maintain the loading capacity of the floor panel.

Data Center Safety

Safety must be a primary consideration in planning installation of VLE system equipment, and is reflected in such choices as where equipment will be located, the rating and capability of electrical, HVAC, and fire-prevention systems that support the operating environment, and the level of personnel training. Requirements of local authorities and insurance carriers will drive decisions as to what constitutes appropriate safety levels in a given environment.

Occupancy levels, property values, business interruption potential, and fire-protection system operating and maintenance costs should also be evaluated. The *Standard for the Protection of Electronic Computer / Data Processing Equipment (NFPA 75)*, the *National Electrical Code (NFPA 70)*, and local and national codes and regulations can be referenced to address these issues.

Emergency Power Control

The data center should be equipped with readily-accessible emergency power-off switches to allow immediate disconnection of electrical power from VLE system equipment. One switch should be installed near each principal exit door so the power-off system can be quickly activated in an emergency. Consult local and national codes to determine requirements for power disconnection systems.

Fire Prevention

The following fire-prevention guidelines should be considered in the construction, maintenance, and use of a data center:

- Store gases and other explosives away from the data center environment.
- Ensure data center walls, floors, and ceilings are fireproof and waterproof.
- Install smoke alarms and fire suppression systems as required by local or national codes, and perform all scheduled maintenance on the systems.

Note – Halon 1301 is the extinguishing agent most commonly used for data center fire suppression systems. The agent is stored as a liquid and is discharged as a colorless, odorless, electrically nonconductive vapor. It can be safely discharged in occupied areas without harm to personnel. Additionally, it leaves no residue, and has not been found to cause damage to computer storage media.

- Install only shatterproof windows, in code-compliant walls and doors.
- Install carbon dioxide fire extinguishers for electrical fires and pressurized water extinguishers for ordinary combustible materials.
- Provide flame-suppressant trash containers, and train personnel to discard combustible waste only into approved containers.
- Observe good housekeeping practices to prevent potential fire hazards.

Site Power Distribution Systems

The following elements of the site power distribution system should be evaluated when planning an installation of VLE system equipment.

System Design

A properly installed power distribution system is required to ensure safe operation of VLE system equipment. Power should be supplied from a feeder separate from one used for lighting, air conditioning, and other electrical systems.

A typical input power configuration, shown in [FIGURE 1-1 on page 19](#), is either a five-wire high-voltage or a four-wire low-voltage type, with three-phase service coming from a service entrance or separately derived source, and with overcurrent protection and suitable grounding. A three-phase, five-wire distribution system provides the greatest configuration flexibility, since it allows power to be provided to both three-phase and single-phase equipment.

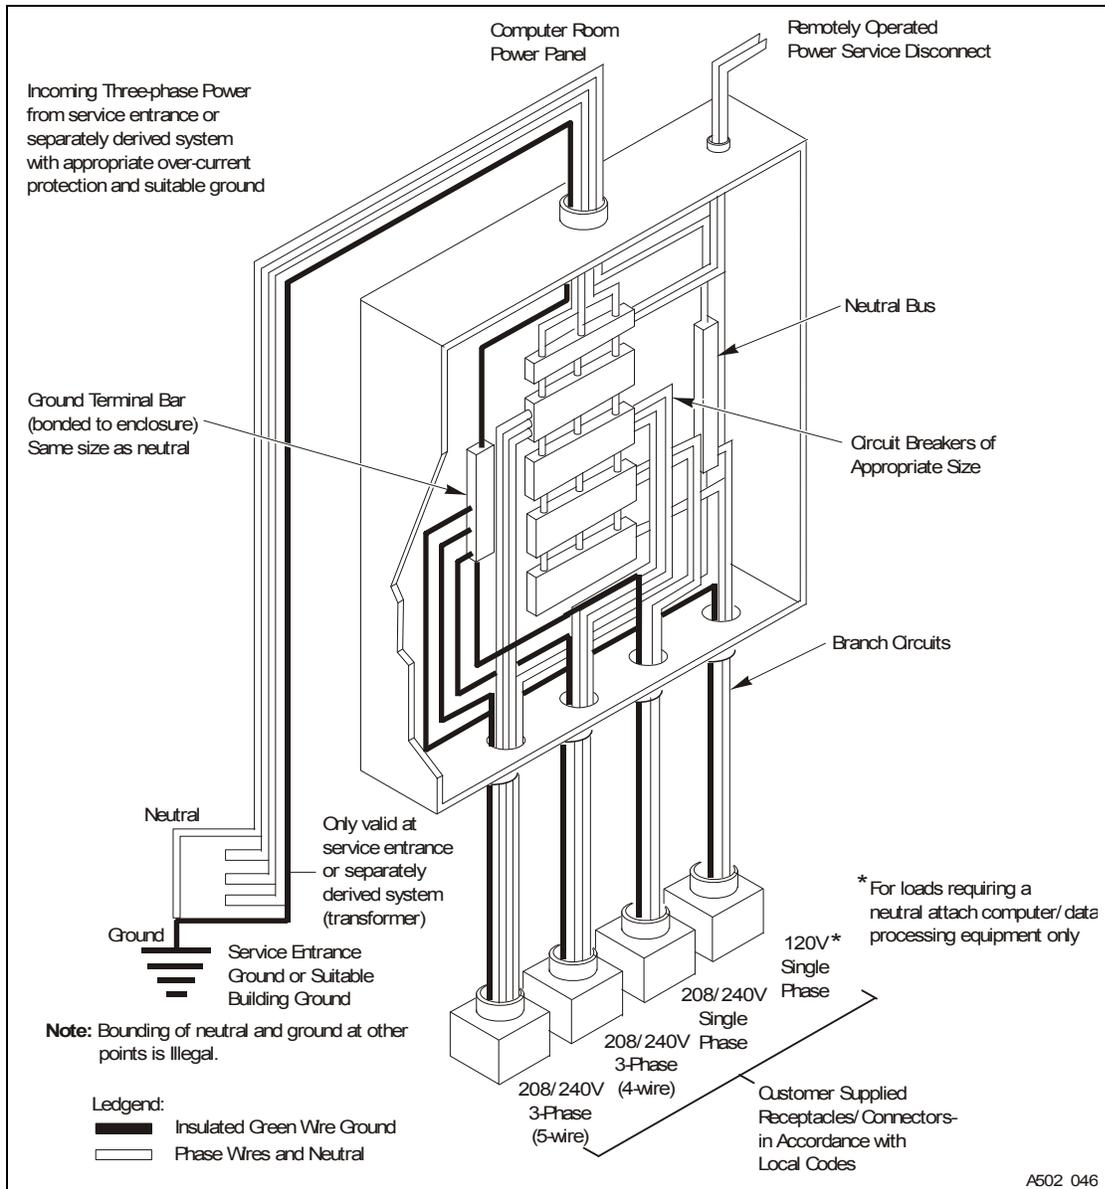


FIGURE 1-1 Site Electrical Power Distribution System

Equipment Grounding

For safety and ESD protection, VLE system equipment must be properly grounded. VLE cabinet power cables contain an insulated green/yellow grounding wire that connects the frame to the ground terminal at the AC source power outlet. A similar insulated green or green/yellow wire ground, of at least the same diameter as the phase wire, is required between the branch circuit panel and the power receptacle that attaches to each cabinet.

Source Power Input

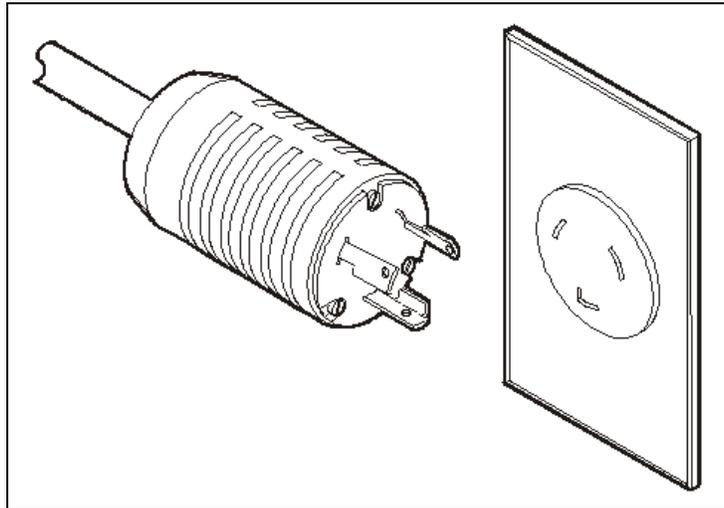
Voltage and frequency ranges at the AC source power receptacle(s) that will supply power to VLE system equipment must be measured and verified to meet the specifications shown in [TABLE 1-3](#).

TABLE 1-3 Source Power Requirements for VLE Equipment

Source Power	Voltage Range	Frequency Range (Hz)
AC, single-phase, 3-wire	170-240	47-63

If you are installing the VLE in the North and South America, Japan and Taiwan, make sure that the designated power sources are NEMA L6-30R receptacles, and make sure that the cabinet power cords are terminated with the required NEMA L6-30P plugs. The factory ships power cords with NEMA L6-30P plugs to North and South America, Japan and Taiwan. Shipments to EMEA and APAC will ship with IEC309 32A 3 PIN 250VAC IP44 plugs.

The figure below shows a NEMA L6-30P plug and L6-30R receptacle.



If you are installing the VLE outside the North and South America, Japan and Taiwan, make sure that designated source-power receptacles meet all applicable local and national electrical code requirements. Then attach the required connectors to the three-wire ends of the cabinet power cords.

Dual Independent Source Power Supplies

VLE cabinets have a redundant power distribution architecture designed to prevent disruption of system operations from single-source power failures. Four 30 Amp power plugs are required.

To ensure continuous operation, all power cables must be connected to separate, independent power sources that are unlikely to fail simultaneously (for example, one to local utility power, the others to an uninterruptible power supply (UPS) system). Connecting multiple power cables to the same power source will not enable this redundant power capability.

Transient Electrical Noise and Power Line Disturbances

Reliable AC source power free from interference or disturbance is required for optimum performance of VLE system equipment. Most utility companies provide power that can properly operate system equipment. However, equipment errors or failures can be caused when outside (radiated or conducted) transient electrical noise signals are superimposed on power provided to equipment.

Additionally, while VLE system equipment is designed to withstand most common types of power line disturbances with little or no effect on operations, extreme power disturbances such as lightning strikes can cause equipment power failures or errors if steps are not taken to mitigate such disturbances.

To mitigate the effects of outside electrical noise signals and power disturbances, data center source power panels should be equipped with a transient grounding plate similar to that shown in [FIGURE 1-2](#).

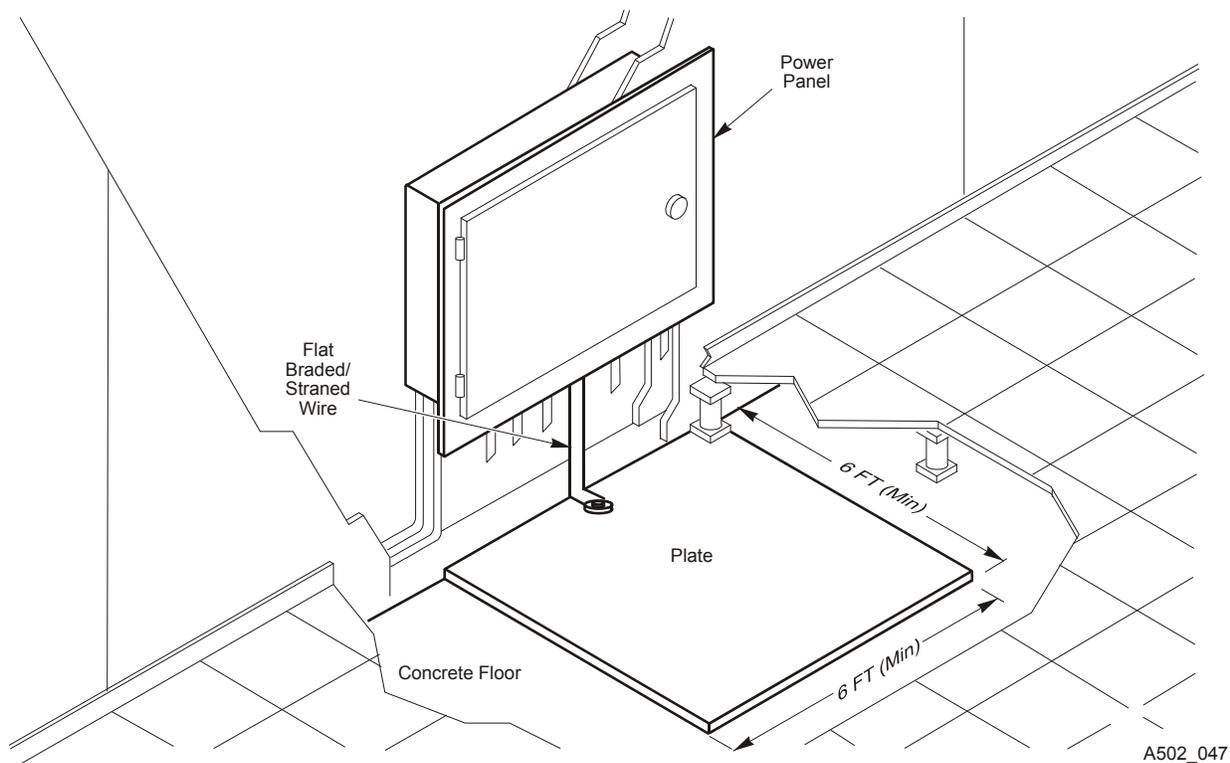


FIGURE 1-2 Transient Electrical Grounding Plate

Electrostatic Discharge

Electrostatic discharge (ESD; static electricity) is caused by movement of people, furniture, and equipment. ESD can damage circuit card components, alter information on magnetic media, and cause other equipment problems. The following steps are recommended to minimize ESD potential in the data center:

- Provide a conductive path from raised floors to ground.
- Use floor panels with nonconducting cores.
- Maintain humidity levels within recommended control parameters.
- Use grounded anti-static work mats and wrist straps to work on equipment.

HVAC Requirements

Cooling and air-handling systems must have sufficient capacity to remove heat generated by equipment and data center personnel. Raised-floor areas should have positive underfloor air pressure to facilitate airflow. If conditions change within a data center (for example, when new equipment is added or existing equipment is rearranged), airflow checks should be done to verify sufficient airflow.

Environmental Requirements and Hazards

VLE system components are sensitive to corrosion, vibration, and electrical interference in enclosed environments such as data centers. Because of this sensitivity, equipment should not be located near areas where hazardous and/or corrosive materials are manufactured, used, or stored, or in areas with above-average electrical interference or vibration levels.

For best performance, equipment should be operated at nominal environmental conditions. If VLE system equipment must be located in or near adverse environments, additional environmental controls should be considered (and implemented where practicable) to mitigate those factors before installation of the equipment.

VLE Planning

This chapter provides information about VLE planning topics.

Satisfying Mainframe Host Software Requirements

For ELS 7.2, support for VLE 1.3 is included in the base level. For ELS 7.0 and 7.1, get the latest SMP/E receive HOLDDATA and PTFs described in [TABLE 2-4](#) and SMP/E APPLY with GROUPEXTEND.

TABLE 2-4 ELS Supporting PTFs for VLE

ELS 7.0	ELS 7.1
L1H16C1	L1H16J6
L1H1672	L1H1674

Satisfying Network Infrastructure Requirements

If possible, do any configuration of IP addresses, network switch(es) for VLANs or other setup (running cables, and so forth) before the VLE arrives to minimize the installation time. Ensure that the network is ready for connection to the VLE as follows:

- Gigabit Ethernet protocol is required on all network switches and routers that are directly attached to VSM5 IFF cards. The IFF card will only do speed negotiation to the 1 Gb speed.
- Switches and routers should support Jumbo(mtu=9000) packets for best performance. If the network is not capable of handling jumbo frames, turn off this capability at the VTSS.

Note – If jumbo frames are enabled then all switches, hubs, or patch panels (including the VLAN and the port channel) between the VLE and its target component must also have jumbo frames enabled.

- Check that you are using the proper (customer-supplied) 1GigE Ethernet cables:
 - CAT5 cables and below are not acceptable for GigE transmission.
 - CAT5E cable: 90 meters is acceptable if run through a patch panel, 100 meters if straight cable.
 - CAT6 cable: 100 meters is acceptable regardless of patch panel configuration.
- StorageTek recommends that if a switch or router is used in the configuration, at least two switches or routers be part of the configuration at each location so that the loss of one unit will not bring down the whole configuration.
- Only one TCP/IP connection is required between a VTSS and a VLE. However, for redundancy, StorageTek strongly recommends that you have a total of four connections between the VTSS and VLE where the VTSS connections are separate IP addresses. Each TCP/IP connection from a specific VTSS to a specific VLE should be to separate VLE interfaces. If you connect all the VTSS connections to the same VLE interface, you have a single point of failure at the VLE interface.

In a VLE multi-node system, the VTSS connections should be spread evenly across all nodes. For example, in a two-node VLE, the VTSS connections should be two on node 1 and the other two on node 2. On a four-node VLE, 1 VTSS connection to each node is recommended. If a switch is involved between the VTSS and VLE, then it is possible to have all four connections to each node of a four-node VLE. Because each VTSS connection represents four drives total, then there would be one drive from each connection to each node for a total of four drives for each node on a four-node VLE.

IP addresses, however, must **never** be duplicated on separate nodes in the VLE for UUI or VTSS. For example, if you have a UUI connection of 192.168.1.1 going to node 1, then do not make a UUI connection on another node using 192.168.1.1 as the IP address! Additionally, if possible, you should never have two interfaces on the same node within the same subnet when configuring IP addresses.

- Similarly, only one UUI connection is required between a VLE and the host, but two are recommended for redundancy, preferably using two independent network paths. Note that these network paths are separate from the connections to the VTSS. For VLE multi-node configurations, if there are multiple UUI connections, make them from separate nodes in the VLE.

Satisfying Oracle Switch Hardware Requirements

The Oracle switch is required for three node or greater VLEs, and can be used for two-node VLEs.

Note – Your VLE is shipped with the following components:

- The dual-port 10GigE NIC cards shown in [FIGURE 2-7 on page 33](#). These cards have installed the pluggable fiber transceivers.
- Two 25 meter fibre optic cables, used to connect VLEs as described above. If you are not going to configure multi-node systems at present, ensure that those cables remain with the VLE decide to configure multi-node systems later.

Order the switch, part X2074A-R to install the switch in a Sun Rack II cabinet.

For each of the first four VLE Nodes that connect through the switch, order two each of the following:

- SFP #X2129A-N
- The appropriate length of LC/LC optic fibre cable, which must be OM3, 850nm, multi-mode, maximum length, including patch panels, of 35M. Two cables per VLE are required to connect to the switch.

For VLE nodes 5 through 7 in the network, **in addition** to the above, order two total (**not** two per node) of the following:

- X2124A-N QSFP parallel fiber optics short wave transceivers.
- X2127A-10M QSFP optical cable splitters. These cable splitters are a single 10M long cable which splits after 9M into four cables for the final one meter in length. Each of these four cable lengths has an LC connector on the end and for the fifth VLE node, you will be able to plug one of the cable ends directly into the VLE node. The overall length of the cable is 10M so that VLE will need to be closer to the switch than VLE's connected by a 35M (or shorter) cable.

For **each** of 6-7 (or if you need more cable for node 5), ensure that you have two of the following:

- LC/LC couplers 10800160-N Spare: LC DUPLEX COUPLING RECEPTACLES.
- LC/LC fibre optic cables no more than 25 meters in length, OM3, 850nm, multimode. The limit is 25 meters because the cables must connect to the QSFP cable, which is 10M, using the couplers.

Satisfying Serviceability Requirements

The VLE product uses a standard Oracle service strategy common with other Oracle products. Automated Service Response (ASR) is used by the VLE as the outgoing event notification interface to notify Oracle Support that an event has occurred on the VLE and the system may require service. Additionally, in combination with ASR, an outgoing email containing details about an ASR event and a Support File Bundle containing VLE log information necessary to investigate any ASR event will also be sent.

The advantages of ASR functionality are well documented in the ASR FAQ available on the My Oracle Support site (<https://support.oracle.com/CSP/ui/flash.html>) in Knowledge Article Doc ID 1285574.1.

Oracle's expectation is that the VLE will be configured to allow outgoing ASR and email communication with Oracle Support. To support VLE outgoing ASR notifications, the customer will need to supply the information in [TABLE 2-5](#) to the installing Oracle Field Engineer.

TABLE 2-5 CAM Configuration Information

Configuration Value	Example
General Configuration - Site Information	
Company Name	Company Inc
Site Name	Site A
City	AnyTown
General Configuration - Contact Information	
First Name	Joe
Last Name	Companyperson
Contact email	joecompanyperson@company.com
Auto Service Request (ASR) Setup - Oracle Online Account Information	
Customer Oracle CSI Login Name	joecompanyperson@company.com
Customer Oracle CSI Login Password	*****
Auto Service Request (ASR) Setup - Internet Connection Settings (Optional)	
Proxy Host Name	web-proxy.company.com
Proxy Port	8080
Proxy Authentication - User Name	
Proxy Authentication - Password	

Note – In [TABLE 2-5 on page 26](#), some fields are not required if a proxy server is not being used or if it does not require an ID and password. If the customer will not provide the CSI email ID and password, then the customer can enter it directly during the install process. ASR registration takes place during the CAM configuration portion of the VLE install. During this part of the install, the VLE will register itself on the Oracle servers as an ASR qualified product.

The customer is then required to log into My Oracle Support (MOS) and approve the registration of the VLE. Until this approval is completed by the customer, the VLE is not capable of auto-generating cases through MOS.

For email notification of event and log information, the customer must also supply the information in [TABLE 2-6](#). If the email server does not require a user name and password, these fields can remain blank.

TABLE 2-6 Notification Setup - Email Configuration Options / ConfCollectStatus

Configuration Value	Example
Email Configuration - SMTP Server Name	SMTP.company.com
Email Configuration - SMTP Server User Name	
Email Configuration - SMTP Server User Password	
Email Recipients	vle@invisiblestorage.com <i>and others as needed</i>

In cases where outgoing communication steps are not completed at the time of installation or not allowed at all, Oracle's options for timely response to events that require support from the Oracle Service team are greatly reduced. The VLE can be configured to send email containing event and log information directly to a designated customer internal email address. A recipient of this email can then initiate a service request directly with Oracle and forward any emails received from the VLE to Oracle Support. In this case, the customer must supply the email address where VLE emails are sent, where this email address can accept emails of up to 5M.

ASR Configuration

By default the VLE will send ASRs through the igb0 port. The site's mail server will be used to send the ASR alerts and the VLE support file bundles. When configuring CAM to send ASRs, it is necessary to input the customer SunSolve email ID and password. When configuring CAM, the customer provides the Oracle CSI email address and password or inputs this information directly into the CAM GUI at the time the CAM Configuration Procedure is performed.

Determining VLE Configuration Values

The following sections tell how to determine configuration values for the VLE.

Note – As noted in the following sections, several software configuration values must match values initially set during configuration of the VLE. Use the `IP_and_VMVC_Configuration.xls` worksheet to record these values so you can pass them on to the personnel who will configure the VLE and the host software.

Determining Values for the Configuration Scripts

To configure the network for VLE, you run the `setup_network` script on each node in a multi-node system (or the only node in a single-node system). If the VLE Name has not been changed from the default, `setup_network` then calls `setup_vle_node`, which configures that node; see [FIGURE 2-3](#).

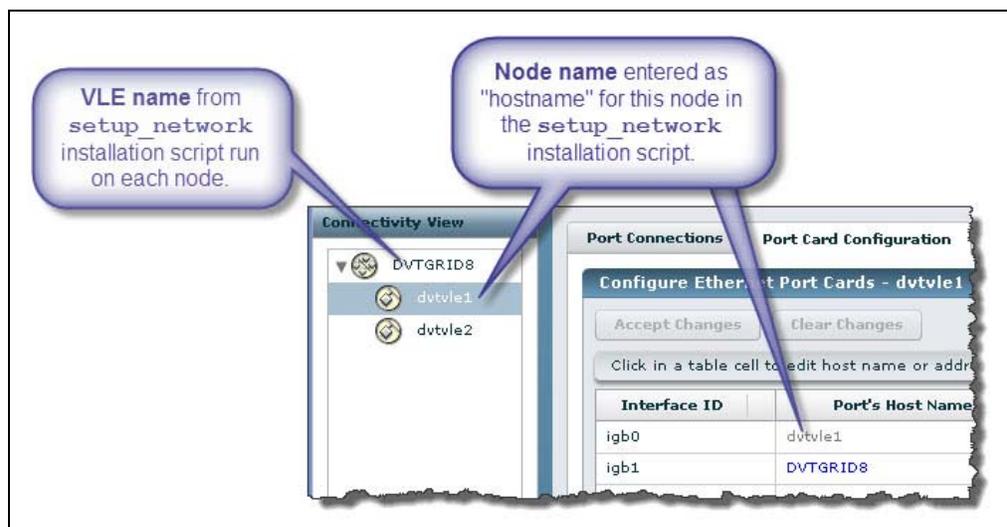


FIGURE 2-3 VLE Name, VLE Number and Node Name

VLE Name and VLE Number

Each VLE node (connected through the same internal network) has a common VLE name and VLE number. The VLE name and number **must be the same** on each node in a multi-node VLE.

The VLE Name must be unique and should **not** be the hostname of any of the servers. The default VLE Name is `VLE-NAME`. You can reset the VLE Name when you run the `setup_vle_node` script. The value must be 1 to 8 characters in length, alphanumeric, uppercase. The name can contain a - (dash) but not at the beginning or the end.

Valid values for the VLE-number are 1-9

In [FIGURE 2-3](#), the VLE-Name and VLE-Number combination is DVTGRID8.

To the host software, the VLE-Name and VLE-Number combination is known as the *subsystem name*, and is specified in the following:

- The STORMNGR parameter value on the VTCS CONFIG TAPEPLEX statement for the TapePlex that connects to the VLE or the NAME parameter on the CONFIG STORMNGR statement (ELS 7.1 and above).
- The STORMNGR parameter value on the VTCS CONFIG RTD statement for the VLE.
- The NAME parameter value on the SMC STORMNGR command that defines the VLE to SMC.
- The STORMNGR parameter value on the SMC SERVER command for the VLE.
- The STORMNGR parameter value on the HSC STORCLAS statement.

Host Name for the Node

As shown in [FIGURE 2-3 on page 28](#), the Host Name for the Node, which is entered on the `setup_network` script, appears as:

- The Port's Host Name for the `igb0` interface ID for the node.
- The Host Name for the node selected in the node navigation tree.

In [FIGURE 2-3 on page 28](#), the Host Name for the node is `dvtvle1`.

Characters can be alpha-numeric (A-Z, a-z, 0-9) or "." or "-". The first and last characters of the string cannot be "." or "-". The name cannot be all-numeric. The name can be up to 512 characters long, though Internet standards and CAM limitations require that the host portion (not including the domain component) be limited to a maximum of 24 characters.

Determining Values for `setup_network`

Required values for the `setup_network` script include the following:

- Hostname for the node; see ["Host Name for the Node" on page 29](#)
- VLE static IP address for port `igb0`
- Network number, which is the base address of the customer subnet
- Netmask
- The default router IP address (Gateway address)
- The network domain name
- The Name Server IP addresses
- Network search names
- NTP server/client setup (server or client, IP addresses of servers) and date/time values

Determining Values for `setup_vle_node`

Required values for the `setup_vle_node` script include the following:

- VLE number and name; see [“VLE Name and VLE Number” on page 28](#).
- Serve Node number (SSN). For multi-node VLEs, each node requires a unique SSN. Valid values for SSN are 1 to 64.
- Server time and date values.

Determining Values for Port Card Configuration

To configure the VLE Ethernet ports, you use the **Connectivity View, Port Card Configuration** tab shown in [FIGURE 2-4](#). The following sections tell how to determine port card configuration values.

The screenshot displays the Oracle Virtual Library Extension GUI. The main window is titled "Configure Ethernet Port Cards - v1e2ga". It features a navigation pane on the left with "Connectivity View" selected, showing a tree structure with "VLE2" and "v1e2ga". The main area has tabs for "Port Connections", "Port Card Configuration", "Port Aggregations", and "Port Status". The "Port Card Configuration" tab is active, showing a table with columns: Interface ID, Port's Host Name, IP Address, Netmask, UUI, and Remote Lis... The table lists several interfaces, with "igb1" selected. Below the table is a "Destination Routes for igb1" pane with buttons for "Add Destination Route", "Get Routes", "Delete Selected Destination Route", "Show ALL Destination Routes", and "Test Connection". A table below this pane shows route details with columns: Type, Gateway Address, Destination IP Address or Name, Remote VLE Name, and Test Result.

Callouts in the image provide the following information:

- Remote VLE Configuration tab removed. Functionality now on Destination Routes pane below.** (Points to the Port Status tab)
- Clear Netmask field by selecting blank item at top of drop-down list.** (Points to the Netmask dropdown menu)
- New "Destination Routes" pane used for defining both Remote VLE connections and Static Routes.** (Points to the Destination Routes pane)
- Type of route defined shown by icons - select "Add Destination Route" button to see descriptions of route types.** (Points to the "Add Destination Route" button)
- Selected Interface.** (Points to the "igb1" row in the table)
- Temporary test button, will be removed.** (Points to the "Test Connection" button)
- Content of bottom pane is filtered by Interface ID selected in top pane. Click this button to show ALL routes for node.** (Points to the "Show ALL Destination Routes" button)

FIGURE 2-4 VLE GUI Port Card Configuration Tab

Interface IDs

The Interface IDs identify the port. You can correlate this ID to ports using the Solaris command `status_vle_ips`. These identifiers are established before the VLE hardware is delivered, and cannot be modified.

VLE Ethernet Ports

On the top of the server, left-hand side at the rear (above the power supplies) there is a sticker that shows these ports with corresponding port IDs as shown in [FIGURE 2-5](#). Note that the sticker shows the location of both the 1GigE ports, shown in [FIGURE 2-5](#), and the 10GigE ports.

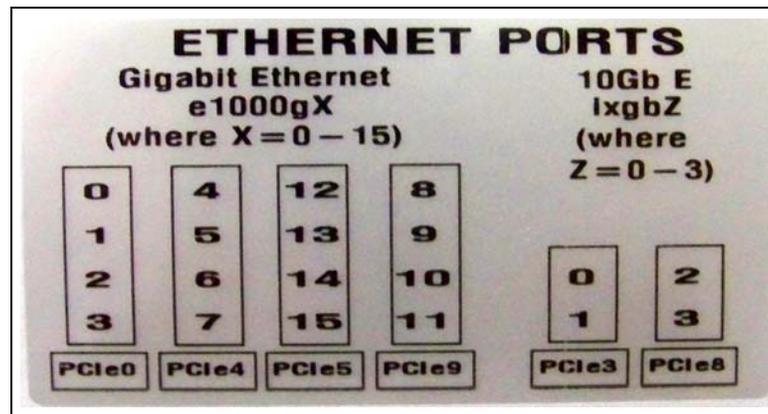


FIGURE 2-5 Ethernet Port Sticker

[FIGURE 2-6](#) shows the 1GigE Ethernet ports in red boxes (e1000g0 to e1000g15) on the rear of the server.

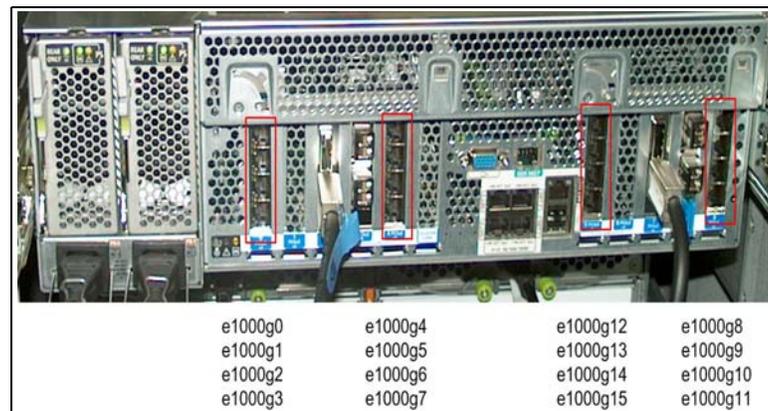


FIGURE 2-6 VLE 1GigE Ethernet Data PortsVLE

The 1GigE Ethernet ports are general purpose ports that can be used for UUI connection, Replication (VLE to VTSS data exchange), Remote Listener (VLE to VLE data exchange) or any combination of all three types.

VLE servers include two dual-port 10GigE NIC cards per server as shown in [FIGURE 2-7](#).

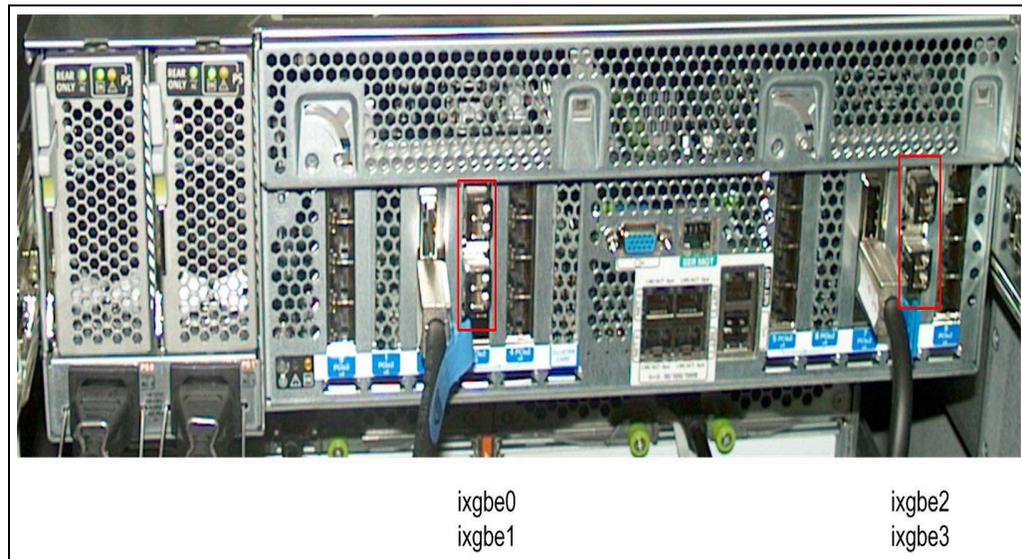


FIGURE 2-7 Dual-port 10GigE NIC cards

As [FIGURE 2-7](#) shows, the 10GigE ports (in the red boxes) have interface IDs of ixgbe0-ixgbe3.

Note – VLEs are shipped with a 1M fiber optic cable connecting ixgb0 to ixgb2. For single node systems, leave the cable connected. VLEs are also shipped with two 25M fiber optic cables, which you can use to connect multi-node system.

Ports ixgbe0 and ixgbe2 are **reserved** for:

- Connections to the Oracle switch for three-node or greater configurations.
- Direct connections to another node in two-node configurations. To connect two nodes, you can do one of the following:
 - Directly connect ixgbe0 on one node to ixgbe0 on the second node, and ixgbe2 on one node to ixgbe2 on the second node.
 - Connect the nodes through the Oracle switch.

Ports ixgbe1 and ixgbe3 are general purpose ports that can be used for UUI connection, Replication (VLE to VTSS data exchange), Remote Listener (VLE to VLE data exchange) or any combination of all three types.

VLE Ethernet Management Ports

The management ports are marked on the back of the case NET0-NET3 as shown in the red box in [FIGURE 2-8](#).

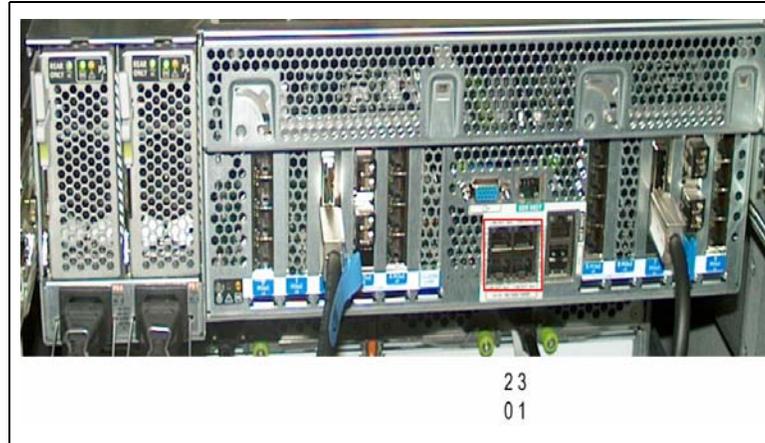


FIGURE 2-8 VLE Ethernet Management Ports

As shown in [FIGURE 2-8](#):

- The management ports (igb0 through igb3) can be on a network segment that is private or public and are typically used as follows:
 - igb0 (NET0) - **Reserved** for connection to the network for ASR traffic and managing the VLE software.
 - igb1 (NET1) - General purpose port, typically used for connection to the network for UUI (control path) traffic.
 - igb2 (NET2) - General purpose port, typically used for redundant UUI connection, or if you want separate ports for separate network segments for the host network and for the sending of ASR alerts.
 - igb3 (NET3) - **Reserved** as dedicated port for service. **Note that** this port can use a single cable and function as **both** the service port and the ILOM port. Do not connect this port to the network. igb3 must remain free and open as an Ethernet port with known access configuration so that it will always be available for service. The preconfigured default IP address for igb3 are:
 - 10.0.0.10 for use as a service port. You use igb3 as a service port to access the VLE CLI.
 - 10.0.0.1 for use as an ILOM port.

Port's Host Name

The value is the machine (host) name for each IP address to be connected to a VTSS or another VLE. Characters can be alpha-numeric (A-Z, a-z, 0-9) or "." or "-". The first and last characters of the string cannot be "." or "-". The name cannot be all-numeric. The name can be up to 512 characters long, though Internet standards and CAM limitations require that the host portion (not including the domain component) be limited to a maximum 24 characters. **Note that** the Port's Host Name for igb0 and igb3 are established during installation, and cannot be changed at the GUI.

IP Address

The IP address assigned to the port, which must be a valid IP v4 address, in the form of "192.68.122.0". Each byte must be 0-255, there must be 4 bytes, numeric only except for the decimal points.

Netmask

The network mask for the port, which must be a valid IP v4 address, in the form of "255.255.255.0". Each byte must be 0-255, there must be 4 bytes, numeric only except for the decimal points.

Replication

Select the check box for each port that will be used for VLE to VTSS data exchange.

UUI

Select the check box for each port that will be used for UUI activity. This port is usually the one used for product configuration and monitoring (including the port used by the GUI browser connection).

Note – Each VLE must have **at least one** UUI connection, and two or more are recommended for redundancy. If you have two or more in a multi-node VLE, spread the UUI connections out over different nodes.

Remote

This check box identifies the port as a “Listener” destination for a VLE-to-VLE data exchange. For VLE to VLE data transfers, the two unused 10GigE connections (ixgbe1 and ixgbe3) or any unused 1GigE connection can be used from any node in a VLE. If each VLE has two or more nodes, StorageTek recommends a **minimum** of one connection from each node to the other VLE. You can run more than one connection from a VLE node to another VLE's node but you should **never** run multiple connections from a VLE node to a single port on the other VLE. If both VLE's have more than one node, StorageTek recommends spreading the VLE to VLE connections across all nodes in each VLE.

For example, VLE1 node 1 has a connection from 192.168.1.1 to VLE2 node 1 at 192.168.1.2. If a second connection is made from VLE node 1, then the connection should **not** go to VLE2 at 192.168.1.2.

For VLE to VLE data transfers, each VLE requires a UUI connection and an VTSS connection. This will ensure VTCS can migrate and recall VTVs from either VLE.

Determining VMVC Range Configuration Values

Ensure that you assign VMVC names and ranges to fit within the site's naming scheme. VMVC names and ranges are set by the CSE during configuration, so it is best to have them assigned before configuration.

As shown in [FIGURE 2-9](#), you use the VLE GUI's **Create New VMVC** dialog box (from the **VMVC View** with a specific node selected in the navigation tree) to specify volser ranges of new VMVCs.

Create New VMVCs - vdev27dr

Enter a 6-character VolSer range for new VMVCs

Incremental:

Prefix: [] Suffix:

First of range

Last of range

Range: VOL00X-VOL20X (21 VMVCs)

OK **Cancel**

FIGURE 2-9 VLE GUI Create New VMVC dialog box

You determine values for each of the fields in [FIGURE 2-9](#) as follows:

- Each of the fields allows 0-6 alpha-numeric characters, with the “assembly” limitations below.
- Alphabetic characters are automatically converted to upper case; leading and trailing spaces in all fields are automatically removed.
- Any of the fields can be empty, allowing the incremental value to be first, last, or in the middle of the volser range name.
- Any of the fields can be either alphabetic or numeric, with field validations to restrict their usage where necessary. For instance, embedded spaces and special characters are not allowed. Invalid field entries are shown with a red box around the field, and selecting the **OK** button will display an error warning.

- The “Incremental” range fields (prefix and suffix) can be either alphabetic or numeric. Field validations ensure that alphabetic and numeric characters are not mixed in either field, the first value must be less than the last value, and max range limits are checked.
- The length of the entire volser name range is constructed by assembly of each field – the length of the prefix + length of ranges + length of the suffix.
For example, you could enter a prefix of AB, a first of range of 001, a last of range of 500 and a suffix of X to build the volser name range of AB001X - AB500X. Similar combinations can be built. But the length of the entire assembly must add up to exactly six characters.
- If the built-up name exceeded the valid 6-character volser name length (like AB0001XY - AB1500XY), clicking the **OK** button displays a warning dialog and does not allow the entry.
- As the range is being built by editing fields, the resulting range is displayed on a line of the dialog just above the **OK** and **Cancel** buttons. The count of VMVCs in the range being built is also displayed in parentheses with the range. If the count exceeds the maximum allowed for the Wildcat box (shown in the “VMVC Counts” fields as Max), the text is displayed in bold orange. At the time the **OK** button is pressed, the current Available count is checked, and if the range exceeds this amount, an error dialog is displayed.
- The suffix string must begin with a different character type (alphabetic vs. numeric) than the incremental range strings. This is for compatibility with VTCS volser name range entry capability. If the range contains the same character type as the beginning of the suffix, the beginning characters of the suffix would be incremented in a range before those in the range fields; that is, VTCS volser name processing is based on character type, not by field-entry of ranges. For example, a GUI entry of 1000 for the First of range, 1094 for the Last of range, and a suffix of 55 would make a range of 100055-109455. On VTCS, this would expand to 100055, 100056, 100057...109455 rather than 100055, 100155, 100255...109455. Because it would be difficult for you to match the latter expansion in VTCS volser name range entry, this construction is prohibited in the GUI.
- If you attempt to define overlapping ranges, only new VMVCs in the range will be added to any already-existing VMVCs (existing VMVCs will not be overwritten or cleared).
- VMVCs have a nominal size of 250 GB (to the host software) and an effective size on the VLE of 1TB (assuming 4:1 compression). [TABLE 2-7](#) shows the maximum VMVCs you can define for each VLE node capacity.

TABLE 2-7 VLE Effective Capacities - Maximum VMVCs Per Node

VLE Effective Capacity	Maximum VMVCs
330 TB	330
660 TB	660
990 TB	990
1320 TB	1320

- The VMVC volser ranges you specify in the VLE GUI must match the volser ranges defined to VTCS!

Planning for Encryption

VLE 1.1 and above provides encryption of VMVCs written to the VLE system. If a VTV is recalled to the VTSS, it is de encrypted at the VLE before recall; therefore, the MVS host software has no knowledge of encryption.

Note –

- The encryption algorithm used is AES-256-CCM. The access key is a 256 bit file.
- FIPS 140-2 certification request has been filed with NIST and is in progress.

Note that encryption is enabled, disabled, and managed at the VLE GUI by a StorageTek CSE or other QSP. Encryption is enabled on a per node basis through an encryption key stored on the node and backed up on a USB device. You can mix encryption and non-encryption nodes in a multi-node VLE because VLE de encrypts VTVs, if required, regardless of where they reside on a multi-node VLE. If, however, you want to encrypt all VTVs on a multi-node VLE, then encryption must be enabled for all nodes.

Some implementation notes:

- **Before** encryption is enabled, there must be **no** VMVCs on the node. Additionally, the USB key backup must be inserted in the node's USB port, and must be writeable and mounted by the OS.
- Similarly, **before** encryption is disabled, recall VTVs that you want to keep to the VTSS, then delete all VMVCs from the node.
- Encryption keys do not expire, so do **not** generate a new key unless you must (for example, to meet security audit requirements). **Before** you assign a new key:
 - The USB key backup must be inserted in the node's USB port, and must be writeable and mounted by the OS.
 - If you are certain you want to generate a new key, ignore the warning and overwrite the old key.

Planning for Deduplication

Deduplication eliminates redundant data in a VLE complex. As the deduplication percentage increases, migration performance can correspondingly improve and network use is reduced.

VLE deduplication is performed at the VLE, so the host job and the VTSS are not affected. When a deduplicated VTV is recalled, the VTV is “rehydrated” (reconstituted) at the VLE before it is recalled to the VTSS. Deduplication occurs on a tape block level within each node and small blocks (less than 4K after compression) are not deduplicated.

Deduplication, which is controlled by the `STORCLAS DEDUP` parameter, increases the effective VLE capacity and is performed by the VLE before the VTV is written to a VMVC. For example, [CODE EXAMPLE 2-1](#) shows deduplication enabled for two Storage Classes.

```
STOR NAME(VLOCAL) STORMNGR(VLESERV1) DEDUP(YES)
STOR NAME(VREMOTE) STORMNGR(VLESERV2) DEDUP(YES)
```

CODE EXAMPLE 2-1 Deduplication Enabled for Local and Remote Storage Classes

The `STORCLAS` statements in [CODE EXAMPLE 2-1](#) specify deduplication for a “local” Storage Class (VLOCAL) on the VLE VLESERV1 and “remote” Storage Class (VREMOTE) on the on the VLE VLESERV2.

[CODE EXAMPLE 2-3](#) shows a Management Class that performs deduplication on the Storage Classes in [CODE EXAMPLE 2-1](#). Any jobs that specify the DEDUP2 Management Class enable deduplication for the referenced Storage Classes.

```
MGMT NAME(DEDUP2) MIGPOL(VLOCAL,VREMOTE)
```

CODE EXAMPLE 2-2 Management Class for Deduplication

Note – Deduplication occurs **only after** the `DEDUP (YES)` policy is set; that is, there is no retroactive deduplication.

Deduplication Guidelines

That’s a quick “how to” for deduplication, now what are the guidelines for what data should and should not be deduplicated? Many sources of mainframe data do **not** benefit from deduplication, such as syslogs. Generally data streams that contain timestamps (where every record is different) will not benefit from deduplication. Backup data streams (where the same records may be written multiple times) will typically benefit from deduplication.

Using the SCRPT Report

After deduplication is enabled, how do you know how well it is working? You can monitor the results with the SCRPT report, as shown in the example in [FIGURE 2-10](#)

Storage Class	STORMNGR	Node	Total MVCs	Capacity (GB)	Used (GB)	Compressed (GB)	Uncompressed (GB)	Reduction Ratio
PROD1	VLELIB1	0	4	1000	200	800	3200	16.0:1
		1	3	750	200	400	1600	8.0:1
		2	5	1250	200	400	1600	8.0:1
		3	4	1000	0	0	0	1.0:1
	VLELIB1	16	4000	600	1600	6400	10.7:1	
Total=		16	4000	600	1600	6400	10.7:1	
{All}	VLELIB1	0	4	1000	200	800	3200	16.0:1
		1	3	750	200	400	1600	8.0:1
		2	5	1250	200	400	1600	8.0:1
		3	4	1000	0	0	0	1.0:1
	VLELIB1	16	4000	600	1600	6400	10.7:1	
Total=		16	4000	600	1600	6400	10.7:1	

FIGURE 2-10 SCRPT Report

In [FIGURE 2-10](#), the approximate reduction ratio for the data, which is Uncompressed Gb divided by Used Gb. The Reduction Ratio, therefore, includes both VTSS compression and VLE deduplication. A larger reduction ratio indicates more effective compression and deduplication.

For example, the VTSS receives 16 Mb of data, compresses it to 4Mb, and writes the compressed data to a VTV. VLE subsequently deduplicates the VTV to 2Mb and writes it to a VMVC. Thus, the reduction ratio is 16Mb divided by 2Mb or 8.0:1.

Because the calculation is done using Mb, it is possible to see 0Gb in the Used or Uncompressed fields, yet see a reduction ratio other than 1.0:1.

Using the MEDVERIFY Utility

You can run the MEDVERIFY utility to verify that VTV data can be read on VMVCs (ELS 7.1 and VLE 1.2 and above only). For VLE, MEDVERIFY ensures that deduplicated VMVCs can be “rehydrated” (reconstituted) when recalled to the VTSS. MEDVERIFY reports on VMVCs that pass or fail verification and also produces XML output.

For example, to verify VTVs on the VMVCs defined in [CODE EXAMPLE 2-1 on page 39](#), enter:

```
MEDVER STOR (VLOCAL)
MEDVER STOR (VREMOTE)
```

In this example:

- MEDVERIFY selects VMVCs in Storage Classes VLOCAL and VREMOTE.
- MAXMVC defaults to 99.
- CONMVC defaults to 1 so only a single VMVC is processed at a time.
- No time out is specified.

Reduced Replication

VLE 1.3 and above offers *Reduced Replication*, which, through VLE-to-VLE replication, allows VTVs to be copied in deduplicated format. The only data copied is data that did not reside on the destination VLE when the copy began. Reduced replication, therefore, reduces the amount of data copied, which lowers network use and copy times. To optimize Reduced Replication, ensure that deduplication is enabled for **both** the source and target Storage Class. Otherwise:

- If deduplication is enabled for the source but not the destination Storage Class, then VTVs are “hydrated” (reconstituted) before being copied.
- If deduplication is enabled for the destination but not the source Storage Class, then VTVs are deduplicated when received at the destination.

For example, [CODE EXAMPLE 2-3](#) shows a Management Class that performs Reduced Replication using the Storage Classes in [CODE EXAMPLE 2-1 on page 39](#).

MGMT NAME(REDREP) MIGPOL(VLOCAL,VREMOTE)
--

CODE EXAMPLE 2-3 Management Class for Reduced Replication

In [CODE EXAMPLE 2-3](#) both Storage Classes are enabled for deduplication. Because the corresponding VLEs are connected and configured for VLE-to-VLE replication, any jobs that specify the REDREP Management Class produce Reduced Replication.

Planning for Link Aggregation

Link aggregation is available for IP configuration for VLE 1.3. A link aggregation consists of multiple interfaces on a VLE node that are configured together as a single, logical unit and share a common IP address. **FIGURE 2-11** shows the **Connectivity View, Port Aggregations** tab, which you use view the pre-defined “internal” aggregation port (such as AggrNode1) and its associated interfaces. You can also define and modify new custom aggregations using this tab.

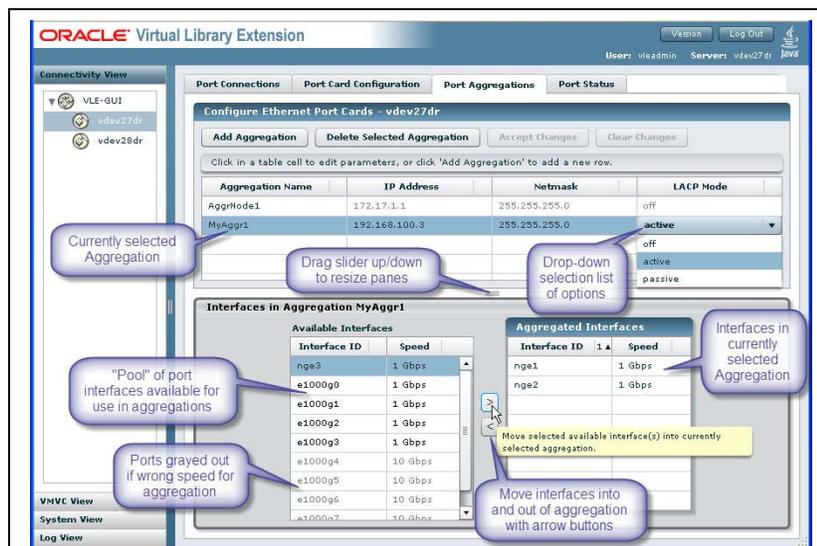


FIGURE 2-11 VLE GUI Connectivity View, Port Aggregations Tab

Benefits of Link Aggregation

Link aggregation provides the following benefits:

- **Less complexity, simpler administration.** Aggregations can simplify VLE configurations by reducing the number of IP addresses required to configure a VLE node, which also prevents drain on the customer address pool. Without link aggregation, more than twenty IP addresses can be required for a fully populated VLE node. Link aggregation can reduce the number of IP addresses to 2, 3, or 4 depending on whether the node has unique Replication, UII, and/or remote VLE IP requirements.
- **Fault tolerance.** With link aggregation, a link can fail and the traffic will switch to the remaining links, thus preventing an outage or job failure.
- **Load balancing and Bandwidth optimization.** The load is balanced by distributing the load of both inbound and outbound traffic across all links in the aggregation. Using all links as one effectively increases bandwidth because traffic is spread evenly across the aggregated links. You can also increase effective bandwidth by increasing the number of links in the aggregation.

For examples of link aggregation, see “[VLE Link Aggregation Examples](#)” on page 61.

Link Aggregation Requirements

- All links in an aggregation must be the same speed. That is, you cannot configure a 1GigE and a 10GigE port in the same aggregation (the VLE GUI does not allow different port speeds in an aggregation).
- The MTU (Maximum Transmission Unit) is configured for the entire aggregation by the Jumbo Frames check-box of the **Port Card Configuration** tab (checking this box sets the MTU (Maximum Transmission Unit) value to 9000 for the aggregation. The switch must support and have the MTU size enabled for all ports within the channel group of the switch.
- An aggregation can consist of a maximum of eight links, which is enforced by the VLE GUI.
- In a switched environment, the first switch from the VLE must support Link Aggregation Control Protocol (LACP) IEEE 802.3ad and be configured for the aggregation mode. The switch is probably a switch in the customer network and is typically administered by a customer network administrator, who will administer the VLE configuration. Ensure that you provide the details of the configuration to the administrator.

Switch Configuration

Note that the terms in the following sections vary between switch vendors. The terms and discussion below are based on CISCO Ethernet switches. Oracle switch terminology is very similar can be found at:

<http://docs.oracle.com/cd/E19934-01/html/E21709/z40016b9165586.html#scrolltoc>

Channel Groups

A channel group is formed in the first switch that is directly connected to the VLE aggregation ports. Other switches or hops in the IP's path need not be aware of the existence of the aggregation. The first switch is responsible for handling the traffic flow to and from the aggregation links. Each channel group is the logical grouping of an aggregation. A channel group is created for each aggregation and contains only the ports of the aggregation. The channel group ties the ports of an aggregation together so the switch can direct traffic to and from the aggregation. Because all ports connected to a channel group are known to be part of the aggregation **do not** connect ports to a channel group that are not part of the aggregation. Each channel group has parameters defined for the type of LACP and so forth, and contains the rules for the aggregation.

VLANs

A typical switch configuration can consist of several VLAN's (Virtual LANs) that connect the VLE to the system components, such as at VTSS or another VLE. A VLAN is a logical grouping of ports in the switch that appear externally as its own isolated switch. The VLAN is typically comprised of one or more channel groups which were created for an aggregation along with the ports of the destination or target components such as the VTSS or another switch in the case of a multi-hop environment.

Jumbo Frames

The MTU (Maximum Transmission Unit) is configured for the entire aggregation by the Jumbo Frames check box of the **Port Card Configuration** tab (checking this box sets the MTU (Maximum Transmission Unit) value to 9000 for the aggregation. If Jumbo frames are enabled then all switches between the VLE and its target components must have jumbo frames enabled as well as for all the ports of the VLAN.

LACP Mode

You can select one of the following LACP modes in the **Aggregation Table** of the **Port Aggregations** tab:

- *Off* - Sometimes referred to as manual mode and indicates LACP datagrams (LACPDUs) are not sent. *Off* is the **only** valid mode without a switch. The non-switched configuration is only valid for VLE to VLE configurations. When using a switch with *Off* mode, LACP is not enabled in the channel group. The switch must be configured to support the aggregation.
- *Passive* - In Passive mode, datagrams are only sent when the switch requests one.
- *Active* - Datagrams are sent to the switch at regular intervals. The timer default of short is used with VLE and is not adjustable with the VLE GUI or CLI.

Policies

P3 is the default VLE policy and is not adjustable through the VLE GUI or CLI.

10GigE Port Aggregations

The 10GigE links can be aggregated for VLE to VTSS, UUI, or VLEto VLE connections. Because UUI traffic is minimal, 10GigE aggregations for UUI only have minimal benefit. 10GigE aggregations that include all three types of connections, however, can prove beneficial. **Note that** for VLE to VTSS configurations, the switch environment typically has both 10GigE and 1GigE connections. In these configurations, the 1GigE VLE ports connect to the switch's 1GigE ports and the VLE 10GigE ports connect to the switch's 10GigE ports. The 10GigE ports would be in a channel group and part of a VLAN that contains both the 1GigE and 10GigE ports.

Monitoring Aggregations

Ensure that you regularly monitor aggregations. If an aggregated link fails, VLE does not generate an ASR because the other links in the aggregation still function, so VLE does not detect the failed link. You cannot monitor the status of the individual links of the aggregation. To display the status of an aggregation, go to the **Connectivity View - Port Status** tab. panel of a VLE node.

Note that if a link goes down, an entry is logged in `/var/adm/messages`. The message file is part of the nightly bundle so the log can be scanned regularly for failed links. The message in the logs looks like the following example:

```
Sep 4 08:30:16 dvtvle3 mac: [ID 486395 kern.info] NOTICE:
e1000g6 link down
```

Types of VLE Aggregations

VLE supports three types of connections, each of which can be aggregated as described in the following sections:

- [“VLE to VTSS Aggregations” on page 45](#)
- [“VLE to VLE Aggregations” on page 45](#)
- [“VLE UII Aggregations” on page 45.](#)

VLE to VTSS Aggregations

Best Practices

- Configure a *minimum* of two aggregations for each VTSS to prevent a total outage if an aggregation fails.
- You can connect multiple VTSSs to the same aggregations. For example, for a VSM5 you can connect IFF0 from each VTSS to one aggregation and connect IFF2 from each VTSS to a second aggregation and so forth. If you are using only two aggregations, then you can connect IFF0 and IFF1 from each VTSS to the first aggregation and so forth.
- Configure links to an aggregation horizontally across the VLE (e1000g0, e1000g4, e1000g8, e1000g12) to prevent an outage to a aggregation if a network adapter fails.

For examples of VLE to VTSS link aggregation, see [“VLE Link Aggregation Examples” on page 61.](#)

VLE to VLE Aggregations

You can aggregate VLE to VLE connections as follows:

- **Non-switched** - In a non-switched configuration, the same interfaces from two VLEs form the connection. The non-switched environment works the same as the internal network of a two-node VLE without a switch. Non-switched environments are limited to point to point configurations only.
- **Switched** - A switched configuration is similar to the configuration described in [“VLE to VTSS Aggregations” on page 45.](#) A channel group is formed in the switch for each aggregation and both channel groups reside in the same VLAN.

With multi-node VLE, a single aggregation from one node can be connected to multiple nodes of another VLE or multiple VLEs in a switched environment.

VLE UII Aggregations

Typically, you use ports `igb1` and `igb2` to make UII connections. In this configuration, aggregate `igb1` and `igb2` to create a *fault-tolerant configuration*: if one of the links fails, the remaining link still provides the UII connection. For additional redundancy on multi-node VLEs, aggregate two UII connections on a second node.

VLE Configuration Examples

This appendix contains the following configuration examples (all examples are direct connect, no switch):

- [“Example 1: One VTSS Connected to One VLE” on page 50](#)
- [“Example 2: Four VTSSs Connected to One VLE” on page 52](#)
- [“Example 3: VLE to VLE Copy” on page 55](#)
- [“Example 4: One VTSS Connected to a 3-Node VLE” on page 58](#)

[“VLE GUI Configuration Screens” on page 48](#) shows the VLE GUI screens used for these examples.

VLE GUI Configuration Screens

The following sections describe the VLE GUI configuration screens used to configure the examples in this appendix. For more information on the VLE GUI, see *Installing, Configuring, and Servicing the VLE Appliance*.

Port Card Configuration Tab

FIGURE A-12 shows the **Port Card Configuration** tab, which is used in all examples. Assume the default **Netmask** value and use whatever value you wish for the **Port's Host Name** field. The table for each example provides the **IP Address** and **Check Box** (Replication, UII, Remote) values that correspond to the **Interface ID** field value.

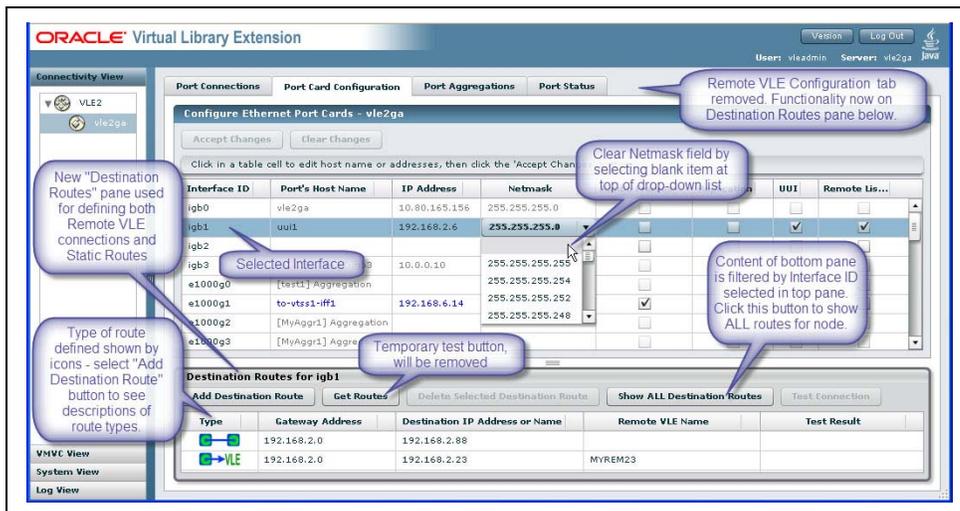


FIGURE A-12 Port Card Configuration Tab

Create New Route/Destination dialog boxes

FIGURE A-13 and FIGURE A-14. You use one of these dialog boxes in “Example 3: VLE to VLE Copy” on page 55.

The screenshot shows the 'Create New Route/Destination' dialog box with the 'Remote VLE Connection' option selected. The 'Interface ID' dropdown is set to 'e100g0'. The 'Gateway Address' field is empty. The 'Remote VLE Node Port Definition' section contains two empty text boxes for 'IP Address or Host Name' and 'Remote VLE Name'. The 'OK' and 'Cancel' buttons are at the bottom right.

FIGURE A-13 Create New Route/Destination dialog box - Remote VLE

The screenshot shows the 'Create New Route/Destination' dialog box with the 'Remote VLE Via Static Route' option selected. The 'Interface ID' dropdown is set to 'igb1' with '192.168.2.6' listed below it. The 'Gateway Address' field contains '192.168.2.0'. A callout bubble points to the 'Remote VLE Via Static Route' option with the text 'Selection modifies available entry fields'. A yellow note box at the bottom left states: 'Note: This operation involves remote VLE configuration, requiring a restart of VLE IP address of interface. Click to copy base to Gateway Address.' The 'OK' and 'Cancel' buttons are at the bottom right.

FIGURE A-14 Create New Route/Destination dialog box - Remote VLE, Static Route

Example 1: One VTSS Connected to One VLE

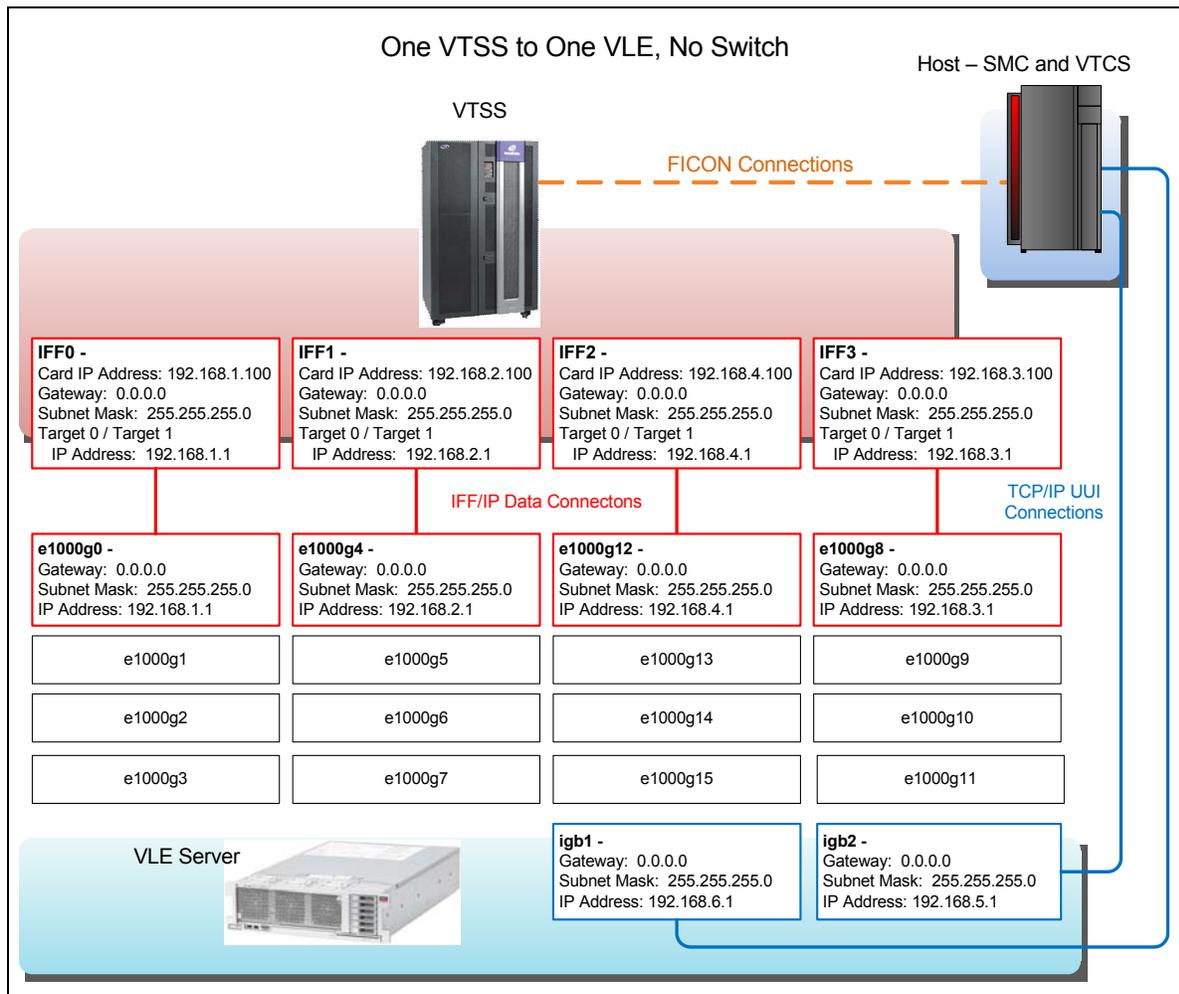


FIGURE A-15 Example 1: One VTSS Connected to One VLE

As [FIGURE A-15 on page 50](#) and [TABLE A-8](#) show, in this one VTSS to one VLE example, two targets on each IFF card connect to a single port on the VLE, where the IP addresses must match. Note that the third octet of the IP addresses is unique to each IFF card to VLE port connection, so these connections share a unique subnet.

Using two targets on each IFF card optimizes performance, because each target represents a socket, which enables a migrate and a recall to occur simultaneously on the same IFF card. Two targets optimizes performance, there is no performance benefit to assigning more than two targets per IFF card to the same VLE port.

TABLE A-8 Example 1 Configuration Values

IFF Card and Target	IPIF Value	Interface ID	IP Address	Check Box
Data Connections				
IFF0 Target 0	0A:0	e1000g0	192.168.1.1	Replication
IFF0 Target 1	0A:1			
IFF1 Target 0	0I:0	e1000g4	192.168.2.1	
IFF1 Target 1	0I:1			
IFF2 Target 0	1A:0	e1000g12	192.168.4.1	
IFF2 Target 1	1A:1			
IFF3 Target 0	1I:0	e1000g8	192.168.3.1	
IFF3 Target 1	1I:1			
UUI Connections				
		igb1	192.168.6.1	UUI
		igb2	192.168.5.1	

Example 2: Four VTSSs Connected to One VLE

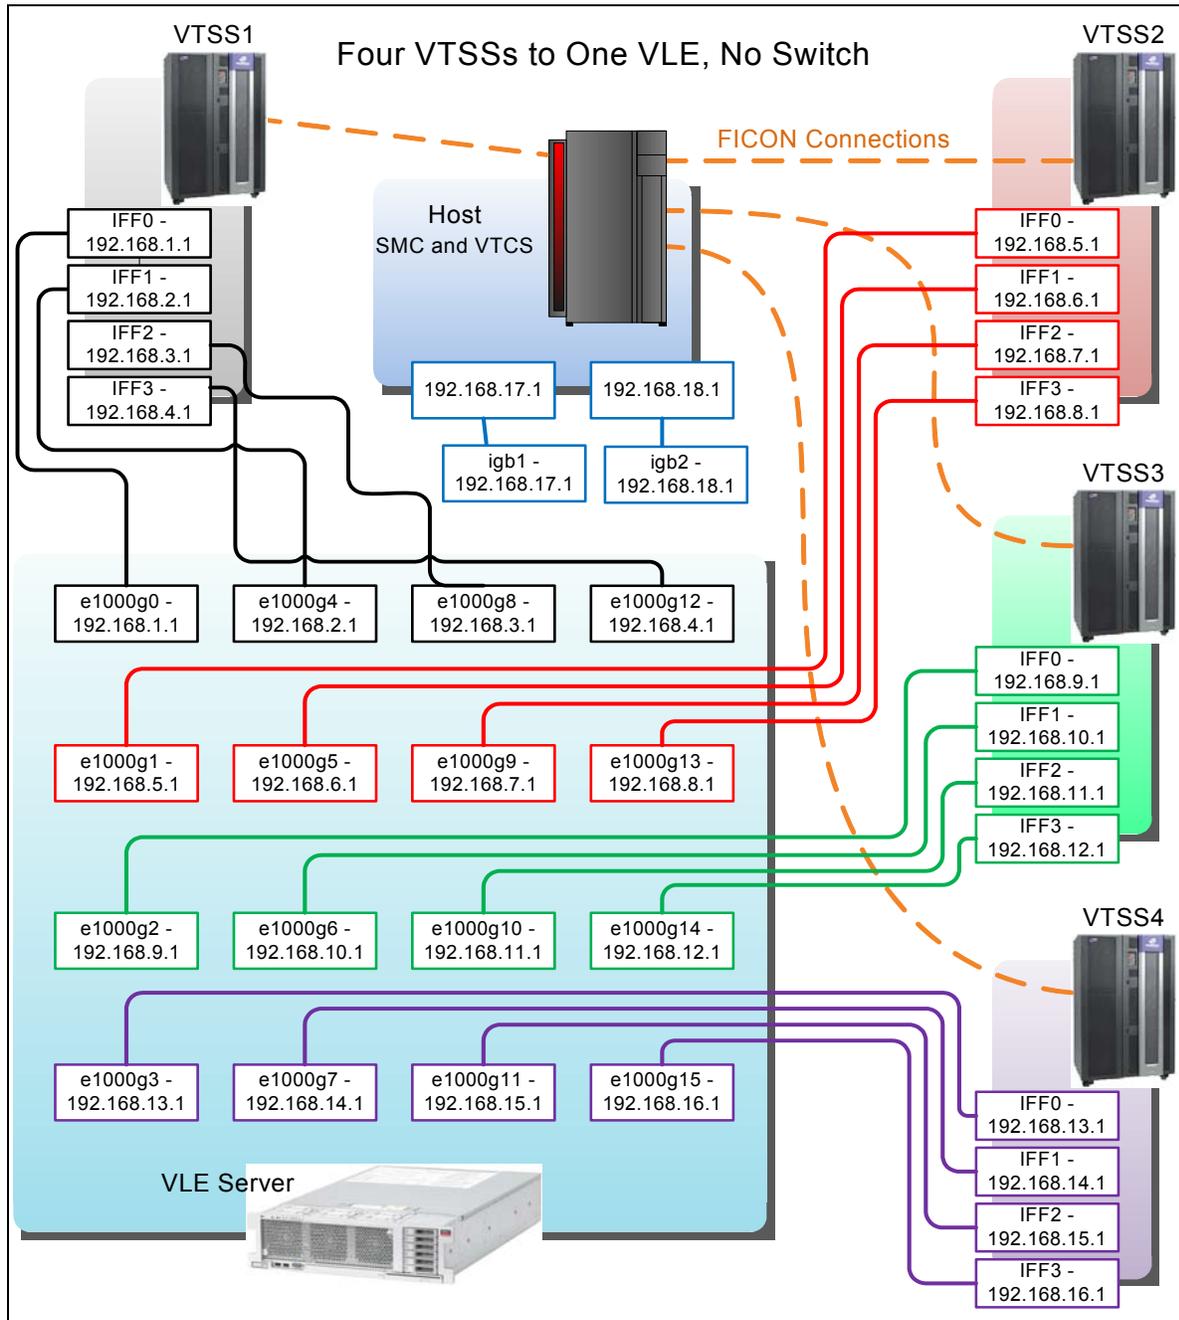


FIGURE A-16 Example 2: Four VTSSs Connected to One VLE

As [FIGURE A-16 on page 52](#) and [TABLE A-9](#) show, in this four VTSS to one VLE example, each IFF target to VLE port connection (where the IP addresses must match) is on its own unique subnet, as shown by the different colors for each subnet (UII connections are shown in [blue](#)).

TABLE A-9 Example 2 Configuration Values

VSM5	IFF Card and Target	IPIF Value	Interface ID	IP Address	Check Box
Data Connections					
VTSS1	IFF0 Target 0	0A:0	e1000g0	192.168.1.1	Replication
	IFF1 Target 0	0I:0	e10004	192.168.2.1	
	IFF2 Target 0	1A:0	e1000g8	192.168.3.1	
	IFF3 Target 0	1I:0	e1000g12	192.168.4.1	
VTSS2	IFF0 Target 0	0A:0	e1000g1	192.168.5.1	
	IFF1 Target 0	0I:0	e1000g5	192.168.6.1	
	IFF2 Target 0	1A:0	e1000g9	192.168.7.1	
	IFF3 Target 0	1I:0	e1000g13	192.168.8.1	
VTSS3	IFF0 Target 0	0A:0	e1000g2	192.168.9.1	
	IFF1 Target 0	0I:0	e1000g6	192.168.10.1	
	IFF2 Target 0	1A:0	e1000g10	192.168.11.1	
	IFF3 Target 0	1I:0	e1000g14	192.168.12.1	
VTSS4	IFF0 Target 0	0A:0	e1000g3	192.168.13.1	
	IFF1 Target 0	0I:0	e1000g7	192.168.14.1	
	IFF2 Target 0	1A:0	e1000g11	192.168.15.1	
	IFF3 Target 0	1I:0	e1000g15	192.168.16.1	

Example 2: Four VTSSs Connected to One VLE

TABLE A-9 Example 2 Configuration Values

UUI Connections					
			igb1	192.168.17.1	UUI
			igb2	192.168.18.1	

Example 3: VLE to VLE Copy

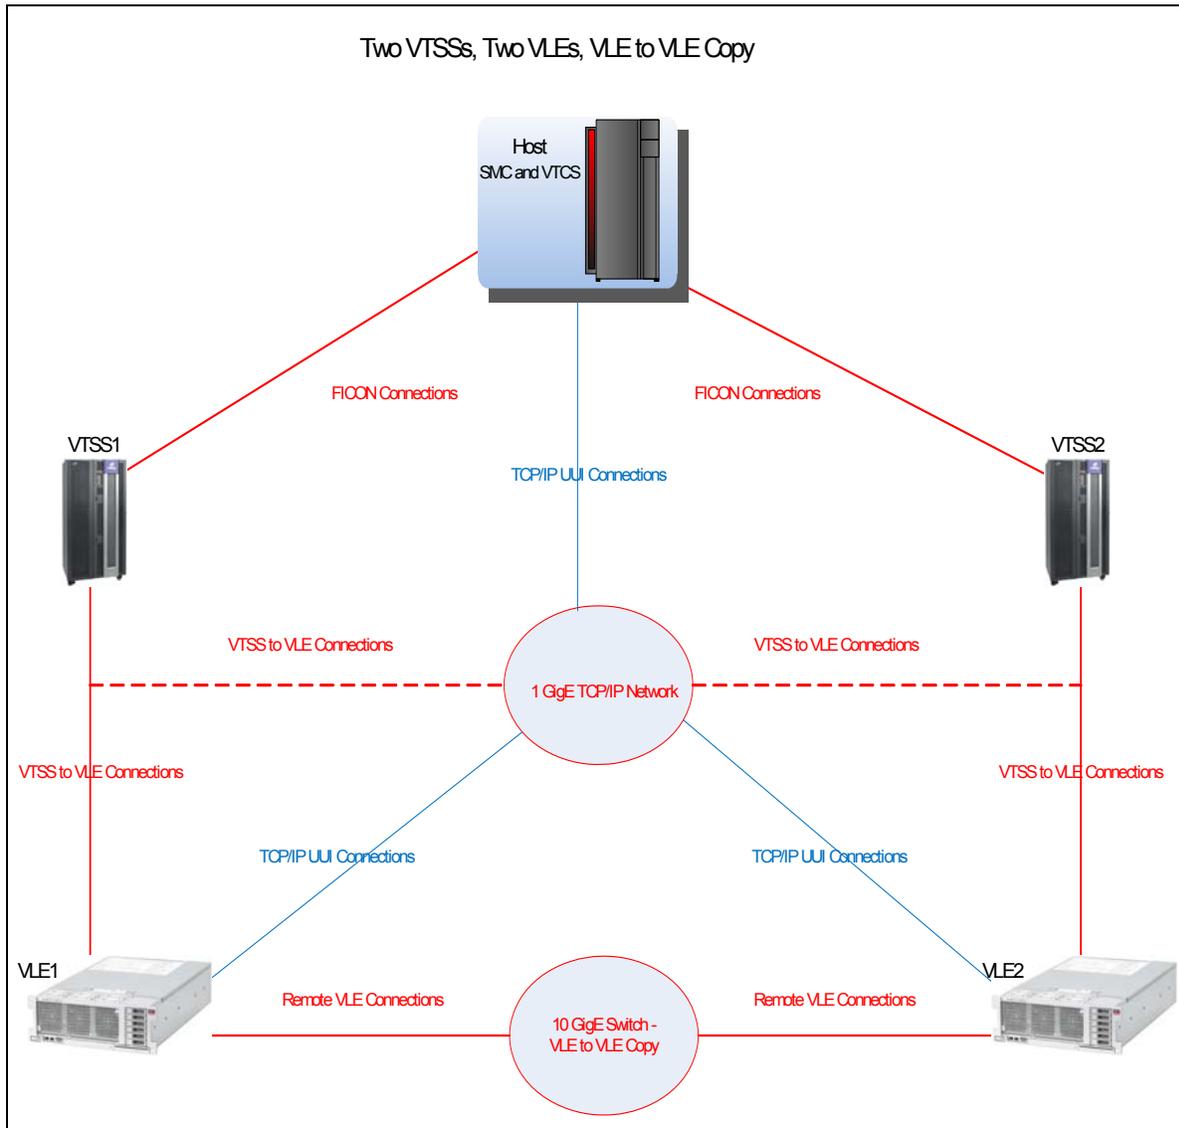


FIGURE A-17 Example 3: VLE to VLE Copy

Example 3: VLE to VLE Copy

As [FIGURE A-17 on page 55](#) and the example values in [TABLE A-10](#) show, in this two VTSSs to two VLEs, VLE to VLE copy example:

- Each VTSS is cross-connected to each VLE for VTSS to VLE copy.
- The VLEs are connected to each other through a network that includes the 10 GigE switch.

Each VTSS, therefore, can migrate a separate VTV copy to each VLE, which provides a redundancy/high availability solution similar to that provided by Clustered VTSS. The default behavior can be that the second copy is made through VLE to VLE connections. To enforce VTSS to VLE migration for the second copy, use the `STORCLAS FROMLST` parameter. For more information, see *Configuring the Host Software for VLE*.

TABLE A-10 Example 3 Configuration Values

VSM5	IFF Card and Target	IPIF Value	VLE, Interface ID	IP Address	Check Box
VLE to VTSS Data Connections					Replication
VTSS1	IFF0 Target 0	0A:0	VLE1, e1000g0	192.168.1.1	
	IFF1 Target 0	0I:0	VLE1, e1000g4	192.168.2.1	
	IFF2 Target 0	1A:0	VLE1, e1000g8	192.168.3.1	
	IFF3 Target 0	1I:0	VLE1, e1000g12	192.168.4.1	
	IFF0 Target 1	0A:1	VLE2, e1000g0	192.168.5.1	
	IFF1 Target 1	0I:1	VLE2, e1000g4	192.168.6.1	
	IFF2 Target 1	1A:1	VLE2, e1000g8	192.168.7.1	
	IFF3 Target 1	1I:1	VLE2, e1000g12	192.168.8.1	

TABLE A-10 Example 3 Configuration Values

VSM5	IFF Card and Target	IPIF Value	VLE, Interface ID	IP Address	Check Box
VLE to VTSS Data Connections					
VTSS2	IFF0 Target 0	0A:0	VLE1, e1000g1	192.168.9.1	Replication
	IFF1 Target 0	0I:0	VLE1, e1000g5	192.168.10.1	
	IFF2 Target 0	1A:0	VLE1, e1000g9	192.168.11.1	
	IFF3 Target 0	1I:0	VLE1, e1000g13	192.168.12.1	
	IFF0 Target 1	0A:1	VLE2, e1000g1	192.168.13.1	
	IFF1 Target 1	0I:1	VLE2, e1000g5	192.168.14.1	
	IFF2 Target 1	1A:1	VLE2, e1000g9	192.168.15.1	
	IFF3 Target 1	1I:1	VLE2, e1000g13	192.168.16.1	
VLE to VLE Data Connections					
			VLE1, ixgbe1	192.168.17.1	Remote
			VLE1, ixgbe3	192.168.18.1	
			VLE2, ixgbe1	192.168.17.2	
			VLE2, ixgbe3	192.168.18.2	
UUI Connections					
			VLE1, igb1	192.168.19.1	UUI
			VLE1, igb2	192.168.20.1	
			VLE2, igb1	192.168.21.1	
			VLE2, igb2	192.168.22.1	

Example 4: One VTSS Connected to a 3-Node VLE

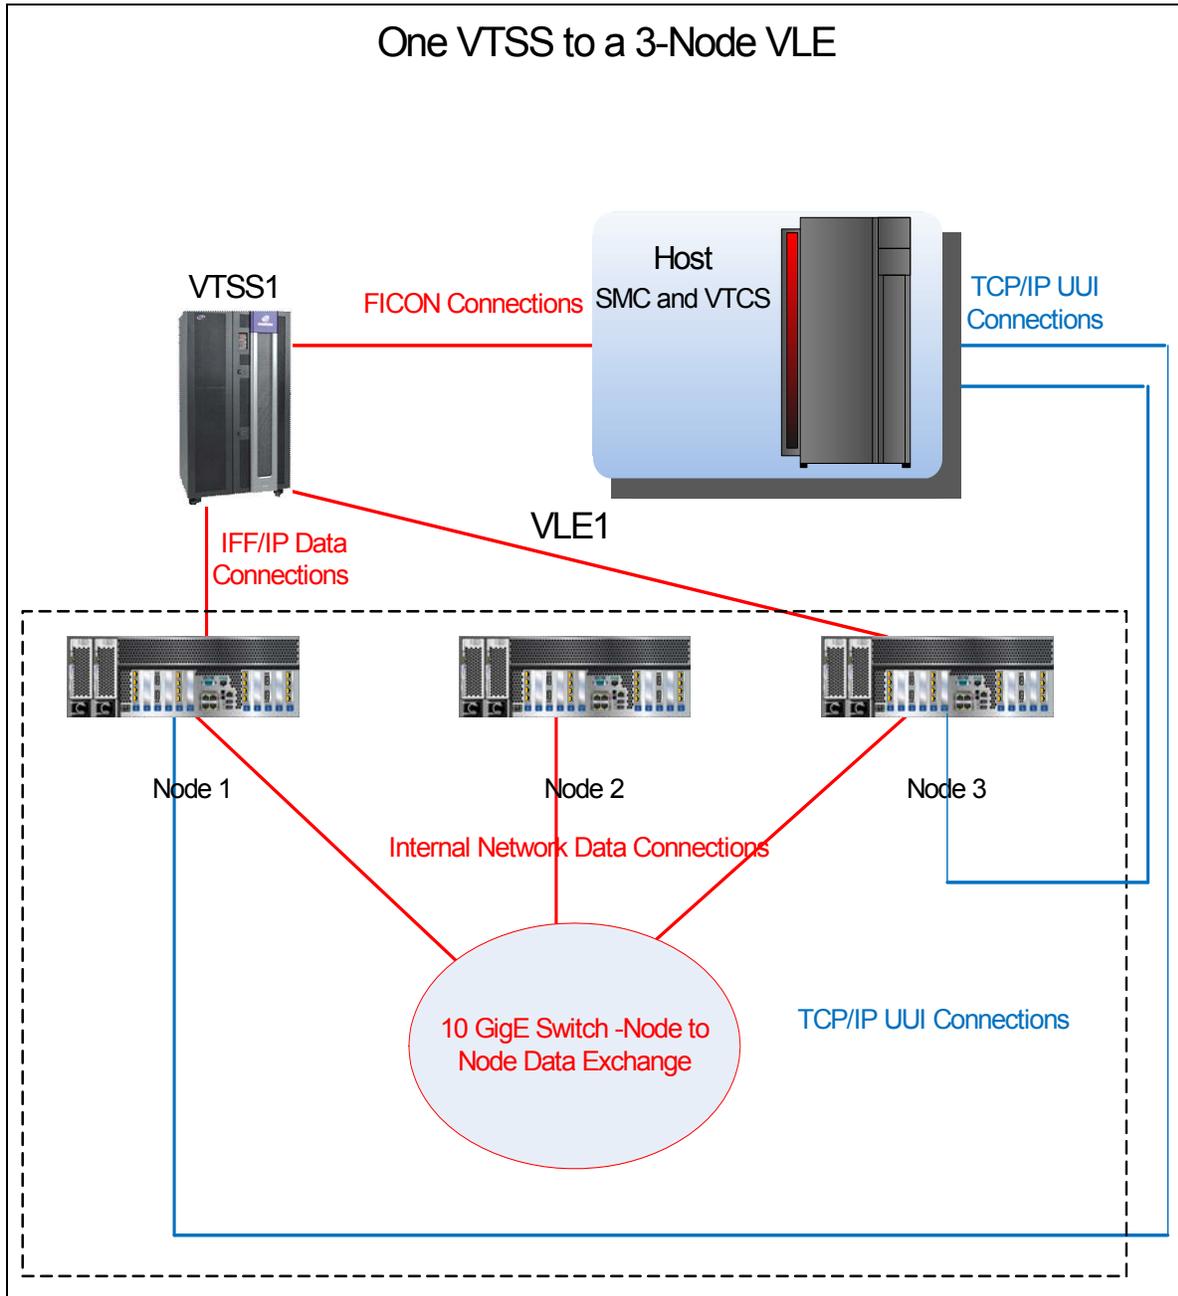


FIGURE A-18 Example 4: One VTSS Connected to a Three-Node VLE

As the example values in [TABLE A-11](#) show, in this one VTSS to a 3-Node VLE example:

- The 10 GigE switch provides an internal network for data exchange between the nodes that make up VLE1.
- To provide redundancy, nodes 1 and 3 both have IFF/IP connections for data exchange with the VTSS and TCP/IP UUI connections to the Mainframe host.
- Node vle1b is a data repository and has no Mainframe host or VTSS connections.

TABLE A-11 Example 4 Configuration Values

IFF Card and Target	IPIF Value	VLE Node, Interface ID	IP Address	Check Box
Data Connections				Replication
IFF0 Target 0	0A:0	Node 1, e1000g0	192.168.1.1	
IFF0 Target 1	0A:1			
IFF1 Target 0	0I:0	Node 1, e1000g4	192.168.2.1	
IFF1 Target 1	0I:1			
IFF2 Target 0	1A:0	Node 3, e1000g0	192.168.3.1	
IFF2 Target 1	1A:1			
IFF3 Target 0	1I:0	Node 3, e1000g4	192.168.4.1	
IFF3 Target 1	1I:1			
UUI Connections				
		Node 1, igb1	192.168.5.1	UUI
		Node 3, igb2	192.168.6.1	

Example 4: One VTSS Connected to a 3-Node VLE

VLE Link Aggregation Examples

This appendix contains the following link aggregation examples:

- [“Example 1: Single-Node VLE Connected to Two VSM5s” on page 63](#)
- [“Example 2: Two-Node VLE Connected to Two VTSSs” on page 64](#)

[“VLE GUI Configuration Screens” on page 62](#) shows the VLE GUI screens used for these examples.

VLE GUI Configuration Screens

The following section shows the VLE GUI Configuration Screen used to configure the examples in this appendix. For more information on the VLE GUI, see *Installing, Configuring, and Servicing the VLE Appliance*.

Port Aggregations tab

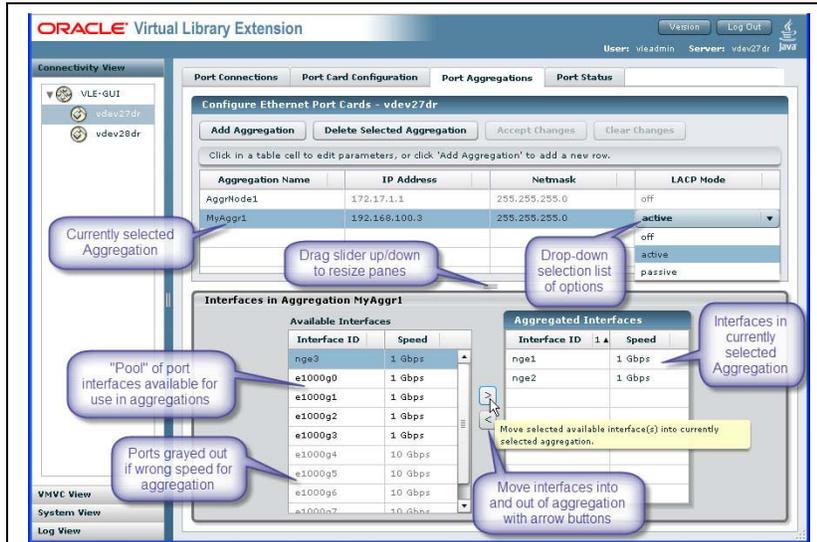


FIGURE B-19 Port Aggregations tab

FIGURE B-19 shows the **Port Aggregations** tab, which is used in all examples.

Example 1: Single-Node VLE Connected to Two VSM5s

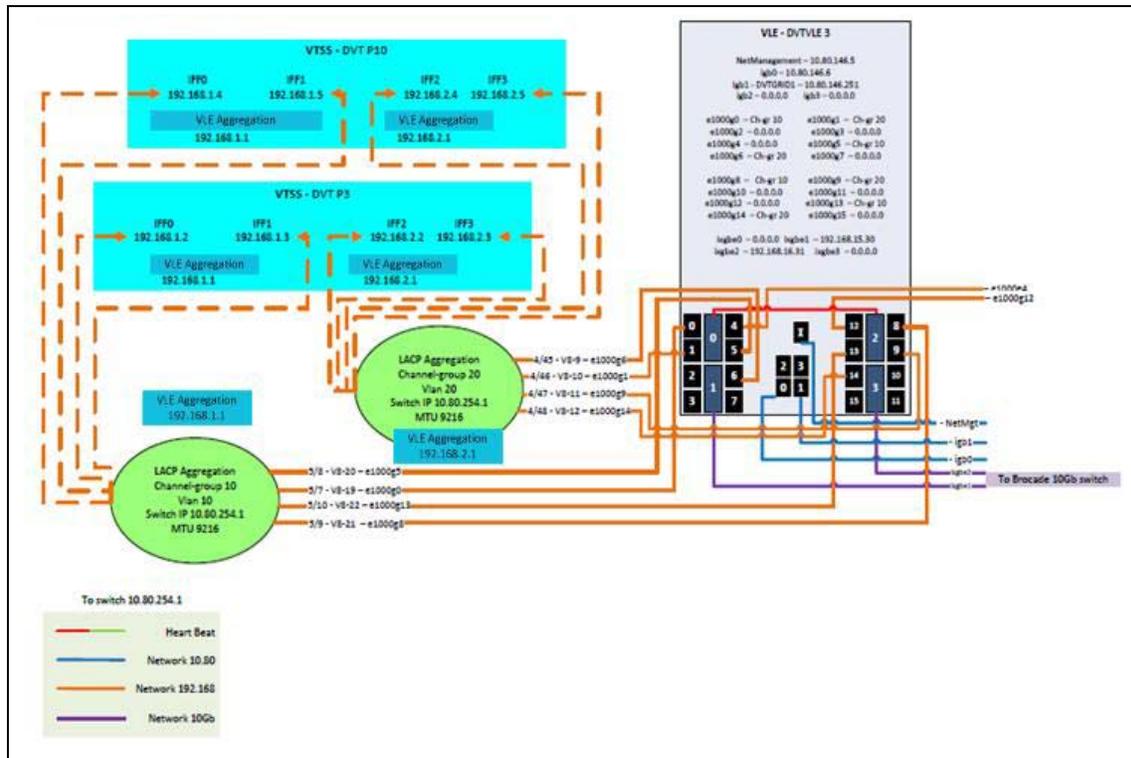


FIGURE B-20 Example 1: Single-Node VLE Connected to Two VTSSs

FIGURE B-20 shows a single-node VLE connected to two VSM5s as follows:

- Ports e1000g0, e1000g5, e1000g8, and g1000g13 are aggregated to the switch's Channel group 10 with an IP address of 192.168.1.1 and the aggregation is connected each VSM5's IFF0 and IFF1 cards.
- Ports e1000g1, e1000g4, e1000g9, and g1000g14 are aggregated to the switch's Channel group 20 with an IP address of 192.168.2.1 and the aggregation is connected each VSM5's IFF2 and IFF3 cards.

Note the following benefits versus a non-aggregated configuration:

- **Simplify the configuration:** Just two IP addresses were required at the VLE (192.168.1.1 and 192.168.2.1).
- **Provide fault tolerance:** The links to each aggregation are configured evenly and horizontally across the VLE to prevent an outage to a aggregation if a network adapter fails. Additionally, both aggregations serve each VSM5, so if an aggregation fails, the other aggregation can continue the connection to the VSM5.
- **Load balancing and bandwidth optimization:** Each aggregation consists of four separate links to distribute the traffic load and optimize bandwidth.

Example 2: Two-Node VLE Connected to Two VTSSs

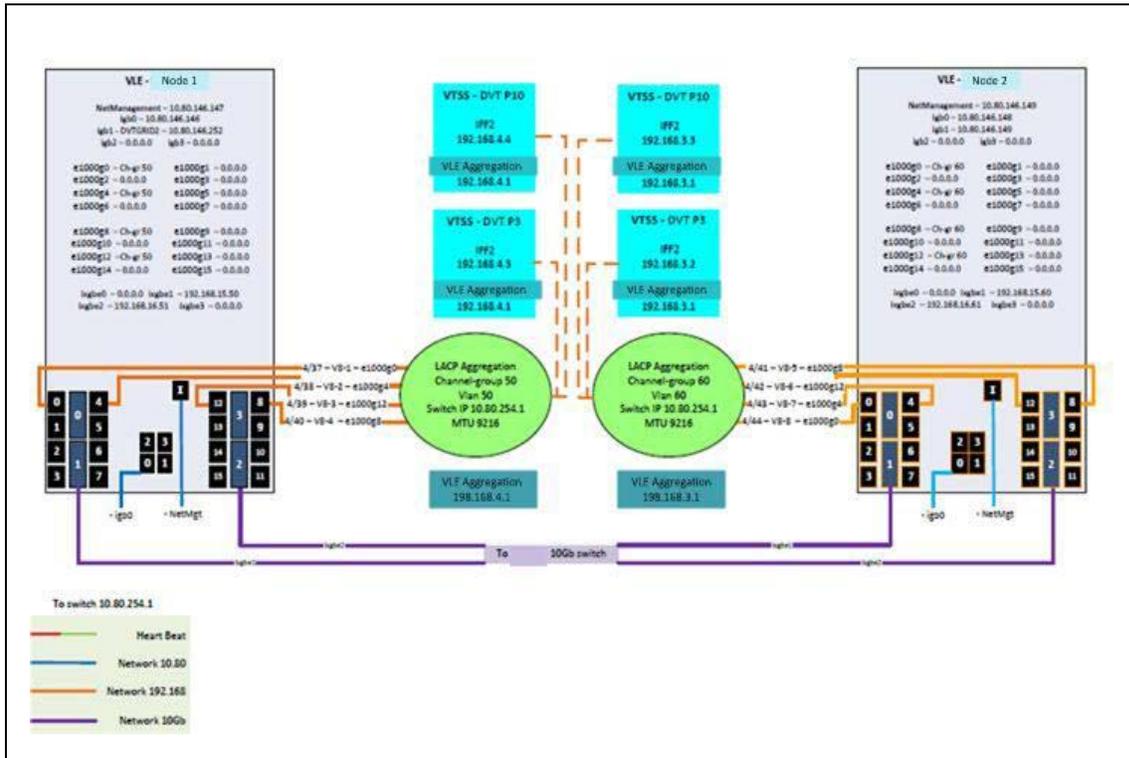


FIGURE B-21 Two-Node VLE Connected to Two VTSSs

FIGURE B-21 shows a two-node VLE connected to two VSM5s as follows:

- Ports e1000g0, e1000g4, e1000g8, and g1000g12 of Node 1 are aggregated to the switch’s Channel group 50 with an IP address of 192.168.4.1 and the aggregation is connected each VSM5.
- Ports e1000g0, e1000g4, e1000g8, and g1000g12 of Node 2 are aggregated to the switch’s Channel group 60 with an IP address of 192.168.3.1 and the aggregation is connected each VSM5.

Note the following benefits versus a non-aggregated configuration:

- **Simplify the configuration:** Just two IP addresses were required at the VLE (192.168.4.1 and 192.168.3.1).
- **Provide fault tolerance:** The links to each aggregation are configured evenly and horizontally across each node to prevent an outage to a aggregation if a network adapter fails. Additionally, both aggregations serve each VSM5, so if an aggregation fails, the other aggregation can continue the connection to the VSM5. Finally, each node services each VSM5, providing additional redundancy.

Load balancing and bandwidth optimization: Each aggregation consists of four separate links to distribute the traffic load and optimize bandwidth.

Controlling Contaminants

Environmental Contaminants

Control over contaminant levels in a computer room is extremely important because tape libraries, tape drives, and tape media are subject to damage from airborne particulates. Most particles smaller than ten microns are not visible to the naked eye under most conditions, but these particles can be the most damaging. As a result, the operating environment must adhere to the following requirements:

- ISO 14644-1 Class 8 Environment.
- The total mass of airborne particulates must be less than or equal to 200 micrograms per cubic meter.
- Severity level G1 per ANSI/ISA 71.04-1985.

Oracle currently requires the ISO 14644-1 standard approved in 1999, but will require any updated standards for ISO 14644-1 as they are approved by the ISO governing body. The ISO 14644-1 standard primarily focuses on the quantity and size of particulates as well as the proper measurement methodology, but does not address the overall mass of the particulates. As a result, the requirement for total mass limitations is also necessary as a computer room or data center could meet the ISO 14644-1 specification, but still damage equipment because of the specific type of particulates in the room. In addition, the ANSI/ISA 71.04-1985 specification addresses gaseous contaminations as some airborne chemicals are more hazardous. All three requirements are consistent with the requirements set by other major tape storage vendors.

Required Air Quality Levels

Particles, gasses and other contaminants may impact the sustained operations of computer hardware. Effects can range from intermittent interference to actual component failures. The computer room must be designed to achieve a high level of cleanliness. Airborne dusts, gasses and vapors must be maintained within defined limits to help minimize their potential impact on the hardware.

Airborne particulate levels must be maintained within the limits of *ISO 14644-1 Class 8 Environment*. This standard defines air quality classes for clean zones based on airborne particulate concentrations. This standard has an order of magnitude less particles than standard air in an office environment. Particles ten microns or smaller are harmful to most data processing hardware because they tend to exist in large

numbers, and can easily circumvent many sensitive components' internal air filtration systems. When computer hardware is exposed to these submicron particles in great numbers they endanger system reliability by posing a threat to moving parts, sensitive contacts and component corrosion.

Excessive concentrations of certain gasses can also accelerate corrosion and cause failure in electronic components. Gaseous contaminants are a particular concern in a computer room both because of the sensitivity of the hardware, and because a proper computer room environment is almost entirely recirculating. Any contaminant threat in the room is compounded by the cyclical nature of the airflow patterns. Levels of exposure that might not be concerning in a well ventilated site repeatedly attack the hardware in a room with recirculating air. The isolation that prevents exposure of the computer room environment to outside influences can also multiply any detrimental influences left unaddressed in the room.

Gasses that are particularly dangerous to electronic components include chlorine compounds, ammonia and its derivatives, oxides of sulfur and petrol hydrocarbons. In the absence of appropriate hardware exposure limits, health exposure limits must be used.

While the following sections will describe some best practices for maintaining an ISO 14644-1 Class 8 Environment in detail, there are some basic precautions that must be adhered to:

- Do not allow food or drink into the area.
- Cardboard, wood, or packing materials must not be stored in the data center clean area.
- Identify a separate area for unpacking new equipment from crates and boxes.
- Do not allow construction or drilling in the data center without first isolating sensitive equipment and any air targeted specifically for the equipment. Construction generates a high level of particulates that exceed ISO 14644-1 Class 8 criteria in a localized area. Dry wall and gypsum are especially damaging to storage equipment.

Contaminant Properties and Sources

Contaminants in the room can take many forms, and can come from numerous sources. Any mechanical process in the room can produce dangerous contaminants or agitate settled contaminants. A particle must meet two basic criteria to be considered a contaminant:

- It must have the physical properties that could potentially cause damage to the hardware.
- It must be able to migrate to areas where it can cause the physical damage.

The only differences between a potential contaminant and an actual contaminant are time and location. Particulate matter is most likely to migrate to areas where it can do damage if it is airborne. For this reason, airborne particulate concentration is a useful measurement in determining the quality of the computer room environment. Depending on local conditions, particles as big as 1,000 microns can become airborne, but their active life is very short, and they are arrested by most filtration devices.

Submicron particulates are much more dangerous to sensitive computer hardware, because they remain airborne for a much longer period of time, and they are more apt to bypass filters.

Operator Activity

Human movement within the computer space is probably the single greatest source of contamination in an otherwise clean computer room. Normal movement can dislodge tissue fragments, such as dander or hair, or fabric fibers from clothing. The opening and closing of drawers or hardware panels or any metal-on-metal activity can produce metal filings. Simply walking across the floor can agitate settled contamination making it airborne and potentially dangerous.

Hardware Movement

Hardware installation or reconfiguration involves a great deal of subfloor activity, and settled contaminants can very easily be disturbed, forcing them to become airborne in the supply air stream to the room's hardware. This is particularly dangerous if the subfloor deck is unsealed. Unsealed concrete sheds fine dust particles into the airstream, and is susceptible to efflorescence -- mineral salts brought to the surface of the deck through evaporation or hydrostatic pressure.

Outside Air

Inadequately filtered air from outside the controlled environment can introduce innumerable contaminants. Post-filtration contamination in duct work can be dislodged by air flow, and introduced into the hardware environment. This is particularly important in a downward-flow air conditioning system in which the sub-floor void is used as a supply air duct. If the structural deck is contaminated, or if the concrete slab is not sealed, fine particulate matter (such as concrete dust or efflorescence) can be carried directly to the room's hardware.

Stored Items

Storage and handling of unused hardware or supplies can also be a source of contamination. Corrugated cardboard boxes or wooden skids shed fibers when moved or handled. Stored items are not only contamination sources; their handling in the computer room controlled areas can agitate settled contamination already in the room.

Outside Influences

A negatively pressurized environment can allow contaminants from adjoining office areas or the exterior of the building to infiltrate the computer room environment through gaps in the doors or penetrations in the walls. Ammonia and phosphates are often associated with agricultural processes, and numerous chemical agents can be produced in manufacturing areas. If such industries are present in the vicinity of the data center facility, chemical filtration may be necessary. Potential impact from automobile emissions, dusts from local quarries or masonry fabrication facilities or sea mists should also be assessed if relevant.

Cleaning Activity

Inappropriate cleaning practices can also degrade the environment. Many chemicals used in normal or “office” cleaning applications can damage sensitive computer equipment. Potentially hazardous chemicals outlined in the “[Cleaning Procedures and Equipment](#)” section should be avoided. Out-gassing from these products or direct contact with hardware components can cause failure. Certain biocide treatments used in building air handlers are also inappropriate for use in computer rooms either because they contain chemicals, that can degrade components, or because they are not designed to be used in the airstream of a re-circulating air system. The use of push mops or inadequately filtered vacuums can also stimulate contamination.

It is essential that steps be taken to prevent air contaminants, such as metal particles, atmospheric dust, solvent vapors, corrosive gasses, soot, airborne fibers or salts from entering or being generated within the computer room environment. In the absence of hardware exposure limits, applicable human exposure limits from OSHA, NIOSH or the ACGIH should be used.

Contaminant Effects

Destructive interactions between airborne particulate and electronic instrumentation can occur in numerous ways. The means of interference depends on the time and location of the critical incident, the physical properties of the contaminant and the environment in which the component is placed.

Physical Interference

Hard particles with a tensile strength at least 10% greater than that of the component material can remove material from the surface of the component by grinding action or embedding. Soft particles will not damage the surface of the component, but can collect in patches that can interfere with proper functioning. If these particles are tacky they can collect other particulate matter. Even very small particles can have an impact if they collect on a tacky surface, or agglomerate as the result of electrostatic charge build-up.

Corrosive Failure

Corrosive failure or contact intermittence due to the intrinsic composition of the particles or due to absorption of water vapor and gaseous contaminants by the particles can also cause failures. The chemical composition of the contaminant can be very important. Salts, for instance, can grow in size by absorbing water vapor from the air (nucleating). If a mineral salts deposit exists in a sensitive location, and the environment is sufficiently moist, it can grow to a size where it can physically interfere with a mechanism, or can cause damage by forming salt solutions.

Shorts

Conductive pathways can arise through the accumulation of particles on circuit boards or other components. Many types of particulate are not inherently conductive, but can absorb significant quantities of water in high-moisture environments. Problems caused by electrically conductive particles can range from intermittent malfunctioning to actual damage to components and operational failures.

Thermal Failure

Premature clogging of filtered devices will cause a restriction in air flow that could induce internal overheating and head crashes. Heavy layers of accumulated dust on hardware components can also form an insulative layer that can lead to heat-related failures.

Room Conditions

All surfaces within the controlled zone of the data center should be maintained at a high level of cleanliness. All surfaces should be periodically cleaned by trained professionals on a regular basis, as outlined in the [“Cleaning Procedures and Equipment”](#) section. Particular attention should be paid to the areas beneath the hardware, and the access floor grid. Contaminants near the air intakes of the hardware can more easily be transferred to areas where they can do damage. Particulate accumulations on the access floor grid can be forced airborne when floor tiles are lifted to gain access to the sub-floor.

The subfloor void in a downward-flow air conditioning system acts as the supply air plenum. This area is pressurized by the air conditioners, and the conditioned air is then introduced into the hardware spaces through perforated floor panels. Thus, all air traveling from the air conditioners to the hardware must first pass through the subfloor void. Inappropriate conditions in the supply air plenum can have a dramatic effect on conditions in the hardware areas.

The subfloor void in a data center is often viewed solely as a convenient place to run cables and pipes. It is important to remember that this is also a duct, and that conditions below the false floor must be maintained at a high level of cleanliness. Contaminant sources can include degrading building materials, operator activity or infiltration from outside the controlled zone. Often particulate deposits are formed where cables or other subfloor items form air dams that allow particulate to settle and accumulate. When these items are moved, the particulate is re-introduced into the supply airstream, where it can be carried directly to hardware.

Damaged or inappropriately protected building materials are often sources of subfloor contamination. Unprotected concrete, masonry block, plaster or gypsum wall-board will deteriorate over time, shedding fine particulate into the air. Corrosion on post-filtration air conditioner surfaces or subfloor items can also be a concern. The subfloor void must be thoroughly and appropriately decontaminated on a regular basis to address these contaminants. Only vacuums equipped with High Efficiency Particulate Air (HEPA) filtration should be used in any decontamination procedure. Inadequately filtered vacuums will not arrest fine particles, passing them through the unit at high speeds, and forcing them airborne.

Unsealed concrete, masonry or other similar materials are subject to continued degradation. The sealants and hardeners normally used during construction are often designed to protect the deck against heavy traffic, or to prepare the deck for the application of flooring materials, and are not meant for the interior surfaces of a supply air plenum. While regular decontaminations will help address loose particulate, the surfaces will still be subject to deterioration over time, or as subfloor activity causes wear. Ideally all of the subfloor surfaces will be appropriately sealed at the time of construction. If this is not the case, special precautions will be necessary to address the surfaces in an on-line room.

It is extremely important that only appropriate materials and methodology are used in the encapsulation process. Inappropriate sealants or procedures can actually degrade the conditions they are meant to improve, impacting hardware operations and reliability. The following precautions should be taken when encapsulating the supply air plenum in an on-line room:

- Manually apply the encapsulant. Spray applications are totally inappropriate in an on-line data center. The spraying process forces the sealant airborne in the supply airstream, and is more likely to encapsulate cables to the deck.
- Use a pigmented encapsulant. The pigmentation makes the encapsulant visible in application, ensuring thorough coverage, and helps in identifying areas that are damaged or exposed over time.
- It must have a high flexibility and low porosity to effectively cover the irregular textures of the subject area, and to minimize moisture migration and water damage.
- The encapsulant must not out-gas any harmful contaminants. Many encapsulants commonly used in industry are highly ammoniated or contain other chemicals that can be harmful to hardware. It is very unlikely that this out-gassing could cause immediate, catastrophic failure, but these chemicals will often contribute to corrosion of contacts, heads or other components.

Effectively encapsulating a subfloor deck in an on-line computer room is a very sensitive and difficult task, but it can be conducted safely if appropriate procedures and materials are used. Avoid using the ceiling void as an open supply or return for the building air system. This area is typically very dirty and difficult to clean. Often the structural surfaces are coated with fibrous fire-proofing, and the ceiling tiles and insulation are also subject to shedding. Even before filtration, this is an unnecessary exposure that can adversely affect environmental conditions in the room. It is also important that the ceiling void does not become pressurized, as this will force dirty air into the computer room. Columns or cable chases with penetrations in both the subfloor and ceiling void can lead to ceiling void pressurization.

Exposure Points

All potential exposure points in the data center should be addressed to minimize potential influences from outside the controlled zone. Positive pressurization of the computer rooms will help limit contaminant infiltration, but it is also important to minimize any breaches in the room perimeter. To ensure the environment is maintained correctly, the following should be considered:

- All doors should fit snugly in their frames.
- Gaskets and sweeps can be used to address any gaps.

- Automatic doors should be avoided in areas where they can be accidentally triggered. An alternate means of control would be to remotely locate a door trigger so that personnel pushing carts can open the doors easily. In highly sensitive areas, or where the data center is exposed to undesirable conditions, it may be advisable to design and install personnel traps. Double sets of doors with a buffer between can help limit direct exposure to outside conditions.
- Seal all penetrations between the data center and adjacent areas.
- Avoid sharing a computer room ceiling or subfloor plenum with loosely controlled adjacent areas.

Filtration

Filtration is an effective means of addressing airborne particulate in a controlled environment. It is important that all air handlers serving the data center are adequately filtered to ensure appropriate conditions are maintained within the room. In-room process cooling is the recommended method of controlling the room environment. The in-room process coolers re-circulate room air. Air from the hardware areas is passed through the units where it is filtered and cooled, and then introduced into the subfloor plenum. The plenum is pressurized, and the conditioned air is forced into the room, through perforated tiles, which then travels back to the air conditioner for reconditioning. The airflow patterns and design associated with a typical computer room air handler have a much higher rate of air change than typical comfort cooling air conditioners so air is filtered much more often than in an office environment. Proper filtration can capture a great deal of particulates. The filters installed in the in-room, re-circulating air conditioners should have a minimum efficiency of 40% (Atmospheric Dust-Spot Efficiency, ASHRAE Standard 52.1). Low-grade pre-filters should be installed to help prolong the life of the more expensive primary filters.

Any air being introduced into the computer room controlled zone, for ventilation or positive pressurization, should first pass through high efficiency filtration. Ideally, air from sources outside the building should be filtered using High Efficiency Particulate Air (HEPA) filtration rated at 99.97% efficiency (DOP Efficiency MILSTD-282) or greater. The expensive high efficiency filters should be protected by multiple layers of pre-filters that are changed on a more frequent basis. Low-grade pre-filters, 20% ASHRAE atmospheric dust-spot efficiency, should be the primary line of defense. The next filter bank should consist of pleated or bag type filters with efficiencies between 60% and 80% ASHRAE atmospheric dust-spot efficiency.

ASHRAE 52-76 Dust spot efficiency %	Fractional Efficiencies %		
	3.0 micron	1.0 micron	0.3 micron
25-30	80	20	<5
60-65	93	50	20
80-85	99	90	50
90	>99	92	60
DOP 95	--	>99	95

Low efficiency filters are almost totally ineffective at removing sub-micron particulates from the air. It is also important that the filters used are properly sized for the air handlers. Gaps around the filter panels can allow air to bypass the filter as it passes through the air conditioner. Any gaps or openings should be filled using appropriate materials, such as stainless steel panels or custom filter assemblies.

Positive Pressurization and Ventilation

A designed introduction of air from outside the computer room system will be necessary in order to accommodate positive pressurization and ventilation requirements. The data center should be designed to achieve positive pressurization in relation to more loosely controlled surrounding areas. Positive pressurization of the more sensitive areas is an effective means of controlling contaminant infiltration through any minor breaches in the room perimeter. Positive pressure systems are designed to apply outward air forces to doorways and other access points within the data processing center in order to minimize contaminant infiltration of the computer room. Only a minimal amount of air should be introduced into the controlled environment. In data centers with multiple rooms, the most sensitive areas should be the most highly pressurized. It is, however, extremely important that the air being used to positively pressurize the room does not adversely affect the environmental conditions in the room. It is essential that any air introduction from outside the computer room is adequately filtered and conditioned to ensure that it is within acceptable parameters. These parameters can be looser than the goal conditions for the room since the air introduction should be minimal. A precise determination of acceptable limits should be based on the amount of air being introduced and the potential impact on the environment of the data center.

Because a closed-loop, re-circulating air conditioning system is used in most data centers, it will be necessary to introduce a minimal amount of air to meet the ventilation requirements of the room occupants. Data center areas normally have a very low human population density; thus the air required for ventilation will be minimal. In most cases, the air needed to achieve positive pressurization will likely exceed that needed to accommodate the room occupants. Normally, outside air quantities of less than 5% make-up air should be sufficient (ASHRAE Handbook: Applications, Chapter 17). A volume of 15 CFM outside air per occupant or workstation should sufficiently accommodate the ventilation needs of the room.

Cleaning Procedures and Equipment

Even a perfectly designed data center requires continued maintenance. Data centers containing design flaws or compromises may require extensive efforts to maintain conditions within desired limits. Hardware performance is an important factor contributing to the need for a high level of cleanliness in the data center.

Operator awareness is another consideration. Maintaining a fairly high level of cleanliness will raise the level of occupant awareness with respect to special requirements and restrictions while in the data center. Occupants or visitors to the data center will hold the controlled environment in high regard and are more likely to act appropriately. Any environment that is maintained to a fairly high level of cleanliness and is kept in a neat and well organized fashion will also command respect from the room's inhabitants and visitors. When potential clients visit the room they will interpret the overall appearance of the room as a reflection of an

overall commitment to excellence and quality. An effective cleaning schedule must consist of specially designed short-term and long-term actions. These can be summarized as follows:

Frequency	Task
Daily Actions	Rubbish removal
Weekly Actions	Access floor maintenance (vacuum and damp mop)
Quarterly Actions	Hardware decontamination Room surface decontamination
Bi-Annual Actions	Subfloor void decontamination Air conditioner decontamination (as necessary)

Daily Tasks

This statement of work focuses on the removal of each day's discarded trash and rubbish from the room. In addition, daily floor vacuuming may be required in Print Rooms or rooms with a considerable amount of operator activity.

Weekly Tasks

This statement of work focuses on the maintenance of the access floor system. During the week, the access floor becomes soiled with dust accumulations and blemishes. The entire access floor should be vacuumed and damp mopped. All vacuums used in the data center, for any purpose, should be equipped with High Efficiency Particulate Air (HEPA) filtration. Inadequately filtered equipment cannot arrest smaller particles, but rather simply agitates them, degrading the environment they were meant to improve. It is also important that mop-heads and dust wipes are of appropriate non-shedding designs.

Cleaning solutions used within the data center must not pose a threat to the hardware. Solutions that could potentially damage hardware include products that are:

- Ammoniated
- Chlorine-based
- Phosphate-based
- Bleach enriched
- Petro-chemical based
- Floor strippers or re-conditioners.

It is also important that the recommended concentrations are used, as even an appropriate agent in an inappropriate concentration can be potentially damaging. The solution should be maintained in good condition throughout the project, and excessive applications should be avoided.

Quarterly Tasks

The quarterly statement of work involves a much more detailed and comprehensive decontamination schedule and should only be conducted by experienced computer room contamination-control professionals. These actions should be performed three to four times per year, based on the levels of activity and contamination present. All room surfaces should be thoroughly decontaminated including cupboards, ledges, racks, shelves and support equipment. High ledges and light fixtures and generally accessible areas should be treated or vacuumed as appropriate. Vertical surfaces including windows, glass partitions, doors, and so on should be thoroughly treated. Special dust cloths that are impregnated with a particle absorbent material are to be used in the surface decontamination process. Do not use generic dust rags or fabric cloths to perform these activities. Do not use any chemicals, waxes or solvents during these activities.

Settled contamination should be removed from all exterior hardware surfaces including horizontal and vertical surfaces. The unit's air inlet and outlet grilles should be treated as well. Do not wipe the unit's control surfaces as these areas can be decontaminated by the use of lightly compressed air. Special care should also be taken when cleaning keyboards and life-safety controls. Specially treated dust wipes should be used to treat all hardware surfaces. Monitors should be treated with optical cleansers and static-free cloths. No Electro-Static Discharge (ESD) dissipative chemicals should be used on the computer hardware, since these agents are caustic and harmful to most sensitive hardware. The computer hardware is sufficiently designed to permit electrostatic dissipation thus no further treatments are required. After all of the hardware and room surfaces have been thoroughly decontaminated, the access floor should be HEPA vacuumed and damp mopped as detailed in the Weekly Actions.

Biennial Tasks

The subfloor void should be decontaminated every 18 months to 24 months based on the conditions of the plenum surfaces and the degree of contaminant accumulation. Over the course of the year, the subfloor void undergoes a considerable amount of activity that creates new contamination accumulations. Although the weekly above floor cleaning activities will greatly reduce the subfloor dust accumulations, a certain amount of surface dirt will migrate into the subfloor void. It is important to maintain the subfloor to a high degree of cleanliness since this area acts as the hardware's supply air plenum. It is best to perform the subfloor decontamination treatment in a short time frame to reduce cross contamination. The personnel performing this operation should be fully trained to assess cable connectivity and priority. Each exposed area of the subfloor void should be individually inspected and assessed for possible cable handling and movement. All twist-in and plug-in connections should be checked and fully engaged before cable movement. All subfloor activities must be conducted with proper consideration for air distribution and floor loading. In an effort to maintain access floor integrity and proper psychrometric conditions, the number of floor tiles removed from the floor system should be carefully managed. In most cases, each work crew should have no more than 24 square feet (six tiles) of open access flooring at any one time. The access floor's supporting grid system should also be thoroughly decontaminated, first by vacuuming the loose debris and then by damp-sponging the accumulated residue. Rubber gaskets, if present, as the metal framework that makes up the grid system should be removed from the grid

work and cleaned with a damp sponge as well. Any unusual conditions, such as damaged floor suspension, floor tiles, cables and surfaces, within the floor void should be noted and reported.

Activity and Processes

Isolation of the data center is an integral factor in maintaining appropriate conditions. All unnecessary activity should be avoided in the data center, and access should be limited to necessary personnel only. Periodic activity, such as tours, should be limited, and traffic should be restricted to away from the hardware so as to avoid accidental contact. All personnel working in the room, including temporary employees and janitorial personnel, should be trained in the most basic sensitivities of the hardware so as to avoid unnecessary exposure. The controlled areas of the data center should be thoroughly isolated from contaminant producing activities. Ideally, print rooms, check sorting rooms, command centers or other areas with high levels of mechanical or human activity should have no direct exposure to the data center. Paths to and from these areas should not necessitate traffic through the main data center areas.

