



Dialogic® PowerVille™ LB
Load Balancer for Real-Time Communications
Technology Guide

Copyright and Legal Notice

Copyright © 2016-2017 Dialogic Corporation. All Rights Reserved. You may not reproduce this document in whole or in part without permission in writing from Dialogic Corporation at the address provided below.

All contents of this document are furnished for informational use only and are subject to change without notice and do not represent a commitment on the part of Dialogic Corporation and its affiliates or subsidiaries ("Dialogic"). Reasonable effort is made to ensure the accuracy of the information contained in the document. However, Dialogic does not warrant the accuracy of this information and cannot accept responsibility for errors, inaccuracies or omissions that may be contained in this document.

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH DIALOGIC® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN A SIGNED AGREEMENT BETWEEN YOU AND DIALOGIC, DIALOGIC ASSUMES NO LIABILITY WHATSOEVER, AND DIALOGIC DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF DIALOGIC PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY INTELLECTUAL PROPERTY RIGHT OF A THIRD PARTY.

Dialogic products are not intended for use in certain safety-affecting situations. Please see <http://www.dialogic.com/company/terms-of-use.aspx> for more details.

Due to differing national regulations and approval requirements, certain Dialogic products may be suitable for use only in specific countries, and thus may not function properly in other countries. You are responsible for ensuring that your use of such products occurs only in the countries where such use is suitable. For information on specific products, contact Dialogic Corporation at the address indicated below or on the web at www.dialogic.com.

It is possible that the use or implementation of any one of the concepts, applications, or ideas described in this document, in marketing collateral produced by or on web pages maintained by Dialogic may infringe one or more patents or other intellectual property rights owned by third parties. Dialogic does not provide any intellectual property licenses with the sale of Dialogic products other than a license to use such product in accordance with intellectual property owned or validly licensed by Dialogic and no such licenses are provided except pursuant to a signed agreement with Dialogic. More detailed information about such intellectual property is available from Dialogic's legal department at 3300 Boulevard de la Côte-Vertu, Suite 112, Montreal, Quebec, Canada H4R 1P8.

Dialogic encourages all users of its products to procure all necessary intellectual property licenses required to implement any concepts or applications and does not condone or encourage any intellectual property infringement and disclaims any responsibility related thereto. These intellectual property licenses may differ from country to country and it is the responsibility of those who develop the concepts or applications to be aware of and comply with different national license requirements.

Dialogic, Dialogic Pro, Veraz, Brooktrout, Diva, BorderNet, PowerMedia, PowerVille, PowerNova, MSaaS, ControlSwitch, I-Gate, Cantata, TruFax, SwitchKit, Eiconcard, NMS Communications, SIPcontrol, Exnet, EXS, Vision, inCloud9, and NaturalAccess, among others as well as related logos, are either registered trademarks or trademarks of Dialogic Corporation and its affiliates or subsidiaries. Dialogic's trademarks may be used publicly only with permission from Dialogic. Such permission may only be granted by Dialogic's legal department at 3300 Boulevard de la Côte-Vertu, Suite 112, Montreal, Quebec, Canada H4R 1P8. Any authorized use of Dialogic's trademarks will be subject to full respect of the trademark guidelines published by Dialogic from time to time and any use of Dialogic's trademarks requires proper acknowledgement.

The names of actual companies and products mentioned herein are the trademarks of their respective owners.

This document discusses one or more open source products, systems and/or releases. Dialogic is not responsible for your decision to use open source in connection with Dialogic products (including without limitation those referred to herein), nor is Dialogic responsible for any present or future effects such usage might have, including without limitation effects on your products, your business, or your intellectual property rights.

Table of Contents

1. Welcome	5
Related Information	5
2. Introduction	6
3. Network Architecture	8
4. PowerVille LB Components	9
Bootstrap Manager Process	10
Service	10
SNMP Module	10
VIP Manager Process and High Availability	10
PowerVille LB Console	12
5. PowerVille LB Functionality	13
6. Message Flows	15
SIP Call Flow	15
HTTP Message Flow	16

Revision History

Revision	Release Date	Notes
3.0	September 2017	Updates to support PowerVille LB 1.5.0.
2.0	November 2016	Updates to support PowerVille LB 1.4.5. PowerVille LB Functionality : Updated section.
1.0	June 2016	Initial release of this document.
Last modified: September 2017		

Refer to www.dialogic.com for product updates and for information about support policies, warranty information, and service offerings.

1. Welcome

This Technology Guide provides an overview and describes the technology for the Dialogic® PowerVille™ LB – Load Balancer for Real-Time Communications (also referred to herein as "PowerVille LB").

Related Information

See the following for additional information:

- PowerVille LB documentation at <http://www.dialogic.com/manuals/lb/powervillelb1-5>.

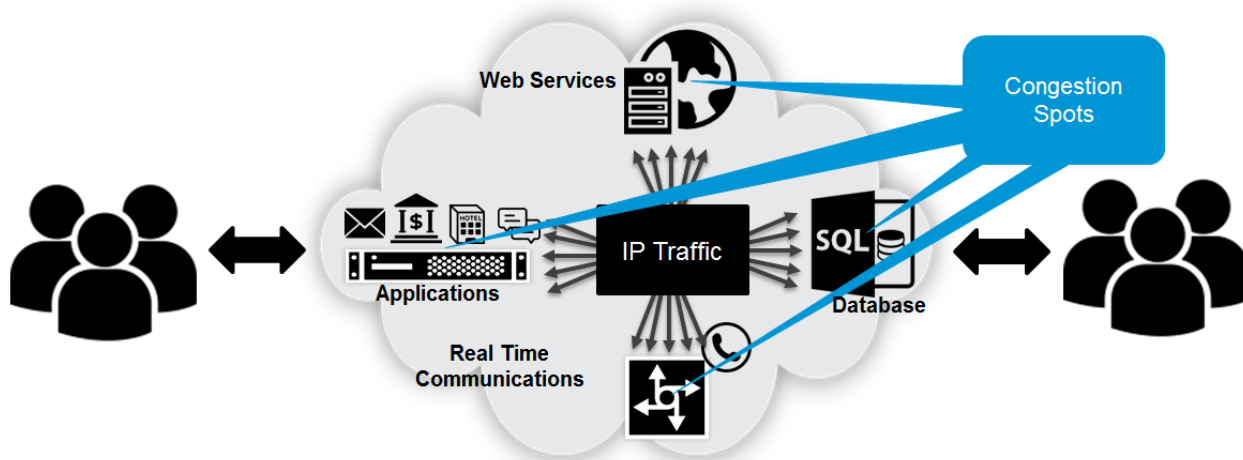
2. Introduction

An Internet Protocol (IP) load balancer provides ultimate control and flexibility of ingress and egress points to deployed, complex solutions. An IP load balancer can present a Virtual IP (VIP) address to the network for single entry point access. The load balancer solution provides numerous features including the ability to distribute incoming network or application traffic across a number of back-end server deployments that require multiple server instances. Ultimately, load balancers are designed to improve service reliability and eliminate capacity bottlenecks for incoming network traffic.

The Dialogic® PowerVille™ LB - Load Balancer for Real-Time Communications ("PowerVille LB") is a software-based solution that provides the following key features:

- **High Performance** - The PowerVille LB is subjected to rigorous performance testing.
- **Cloud Ready** - The PowerVille LB has evolved alongside emerging cloud technology and is continually adapted to meet customer needs.
- **RTC Purpose Built** - The PowerVille LB is fully optimized for network traffic and is uniquely designed to meet challenges for today's demanding Real-Time Communications (RTC) infrastructure in both carrier and enterprise applications.

The following diagram provides examples of where the PowerVille LB can be positioned in solution deployments.



The development of the PowerVille LB has been determined by its usage in Voice over IP (VoIP) solutions, which has resulted in a strong technology focus on servicing RTC and, more specifically, in the area supporting the Session Initiation Protocol (SIP). The PowerVille LB has sophisticated SIP support to maintain high-performance figures while providing "smart technology" such that the most complex deployment scenarios can be met.

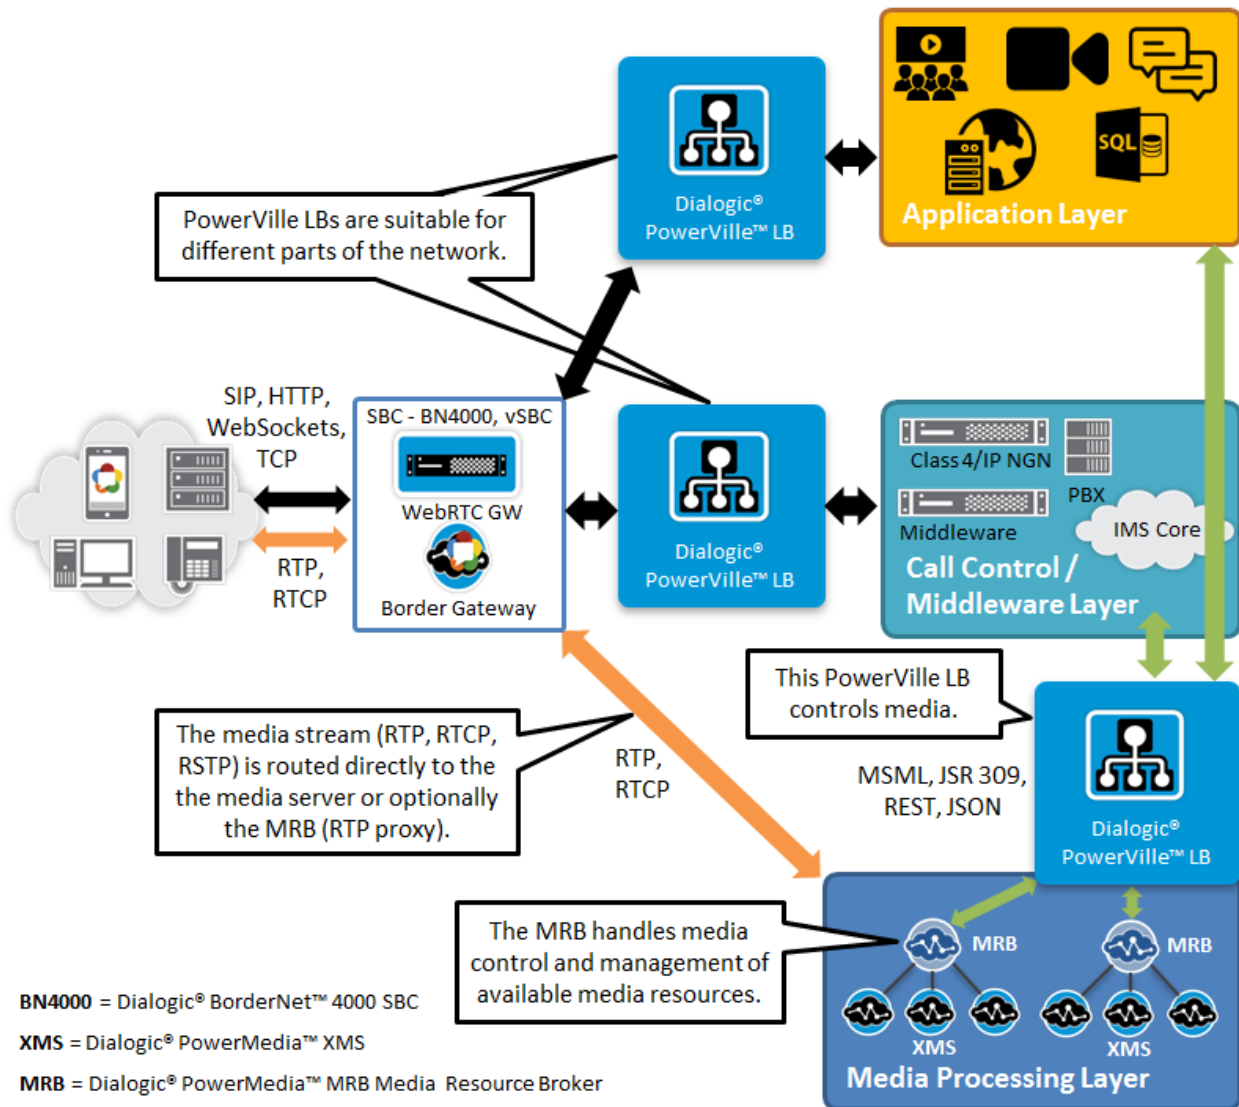
These are some examples of the RTC network requirements that the PowerVille LB meets:

- High-latency sensitivity
- Carrier-class reliability and failover
- Strong performance and throughput
- Cloud readiness
- Service awareness

Application developers, service providers, and enterprises can use the PowerVille LB to dynamically scale, distribute, and manage traffic associated with a diverse set of real-time and non-real-time applications deployed in today's networks across disparate applications and datacenters.

3. Network Architecture

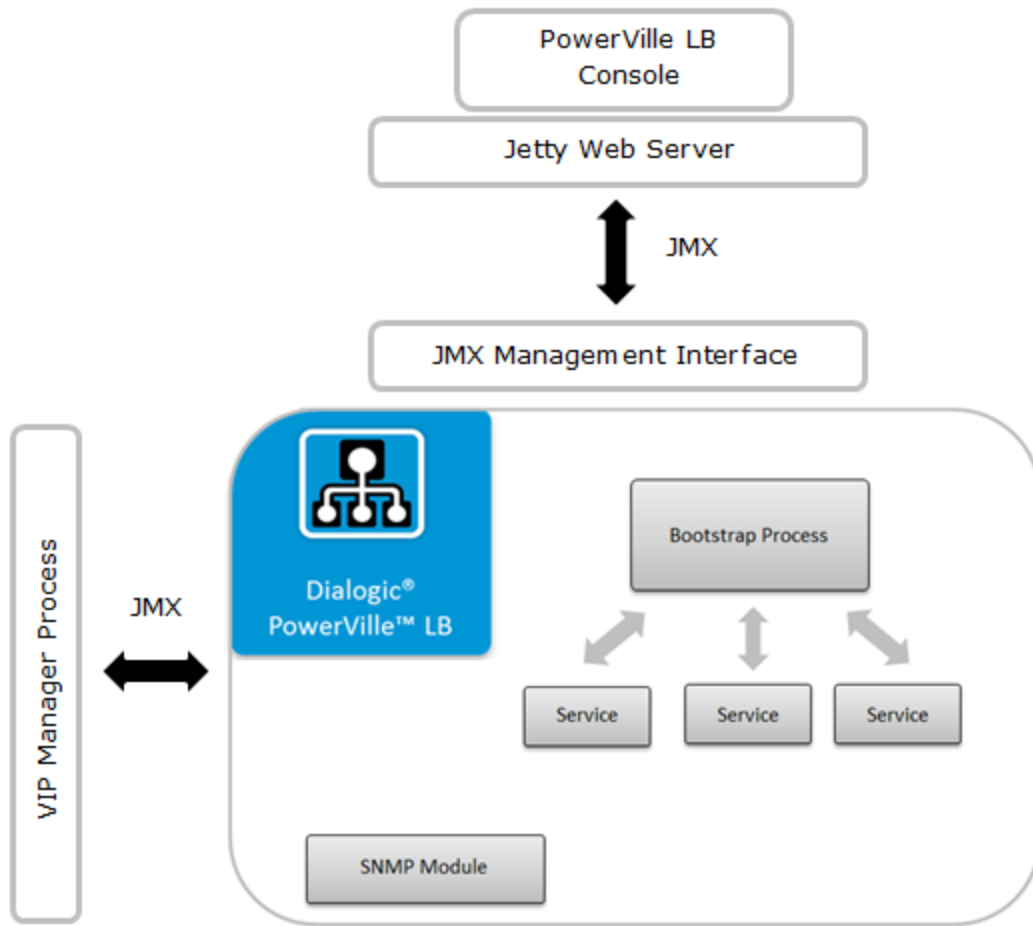
The PowerVille LB provides an extremely flexible and powerful solution that can adapt to requirements demanded of complex network deployments. The ability to mix and match support for numerous protocol interactions enables the PowerVille LB to reside at the heart of complex deployments, which provides an abstracted and simplified offering. In the following high-level example diagram, the PowerVille LB is shown acting at the heart of typical, complex interactions between clients and servers.



In the diagram, interactions between the Application Layer, the Call Control Layer, and the Media Processing Layer illustrate how the PowerVille LB can integrate with other solutions and work with protocols such as SIP, HTTP, WebSockets, etc., which are increasingly used to form modern communications solutions. The PowerVille LB balances ingress traffic from devices such as Session Border Controllers (SBC), Border Gateways, and WebRTC gateways in conjunction with internal core traffic. A single instance of the PowerVille LB can be used to manage multiple services, or multiple instances of the PowerVille LB can be used to present a smaller number of services.

4. PowerVille LB Components

The PowerVille LB has a number of important components. The following diagram provides an overview of these primary components and how they fit together.



The following major components of the PowerVille LB are described in the next sections:

- [Bootstrap Manager Process](#)
- [Service](#)
- [SNMP Module](#)
- [VIP Manager Process and High Availability \(HA\)](#)
- [PowerVille LB Console](#)

Bootstrap Manager Process

The Bootstrap Manager process is the administrative guardian of the PowerVille LB. It is responsible for a number of system-level tasks and monitoring, which include the following:

- The life cycle is managed for all configured PowerVille LB network services. This includes starting, restarting, and stopping PowerVille LB services when required.
- The PowerVille LB services are monitored to detect failure and take action if required.
- The Bootstrap Manager process interacts with the SNMP Module to produce appropriate SNMP traps when configured.
- The Bootstrap Manager process interacts with the VIP Manager process such that VIP addresses are appropriately presented to the network for configured services.

Service

Each service configured on the PowerVille LB constitutes a separate and individual process that is managed by the Bootstrap Manager process. For example, a PowerVille LB instance configured to present a SIP and HTTP service to the network has two discrete PowerVille LB service processes being managed by the Bootstrap Manager process. There are a number of transport protocols for the PowerVille LB service, which are broken down further into stateful or stateless service.

Stateful and stateless services offer different functionality. Unless advised by Dialogic, use the SIP stateless service instead of the SIP stateful service. Stateless service supports a higher rate through the PowerVille LB. In certain testing scenarios, the SIP stateless service was found to be three times faster than the SIP stateful service. The SIP stateless service does not handle 3xx redirects. Instead, the 3xx redirects are forwarded back to the client to handle. Unlike the SIP stateless service, the SIP stateful service does not support PRACKs. Both stateful and stateless SIP services support high availability (HA) failover.

Similarly, the HTTP stateless service cannot do HTTPS offloading and source affinity across multiple connections to the same client, but it can handle much higher throughputs than the HTTP stateful service.

For more information on configuring services, refer to the *Dialogic® PowerVille™ LB – Load Balancer for Real-Time Communications Installation and Operations Guide*.

SNMP Module

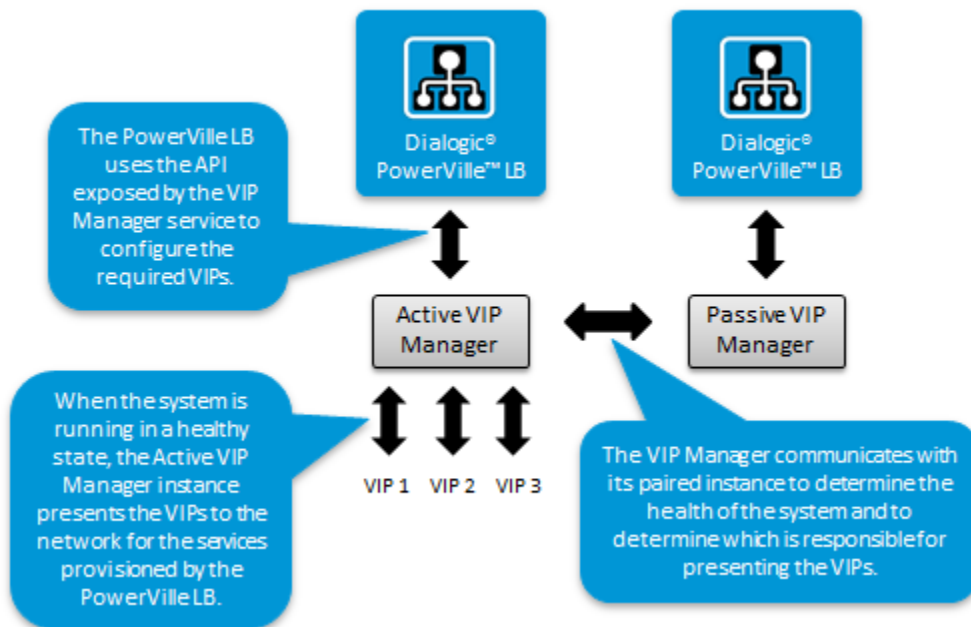
The PowerVille LB SNMP Module interacts with the Bootstrap Manager process to produce appropriate SNMP traps which are then dispatched to configured SNMP trap listeners. The desired trap listeners can be provisioned using the PowerVille LB console. A full list of traps generated by the PowerVille LB can be found in the *Dialogic® PowerVille™ LB – Load Balancer for Real-Time Communications Installation and Operations Guide*.

VIP Manager Process and High Availability

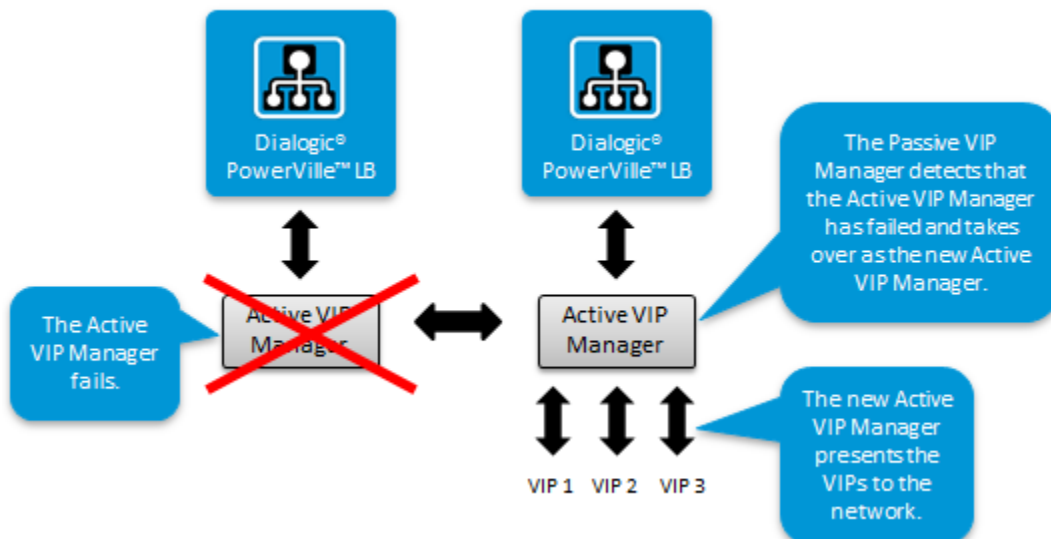
The VIP Manager is a service utilized by the PowerVille LB primarily for presenting appropriate VIP addresses to the network. Interactions take place with the Bootstrap Management process such that PowerVille LB services are appropriately configured. The VIP Manager plays a vital role in supporting High Availability (HA) for services being managed by the PowerVille LB.

When deployed, a pair of PowerVille LB instances are configured to provide HA to the network services being offered. Each instance of the PowerVille LB connects to a local VIP Manager process to configure the VIP addresses that are required to present services to the network. The two VIP Manager instances that comprise a default PowerVille LB deployment collaborate to determine which is responsible for VIP presentation to the network. Should a local failure occur, the surviving VIP Manager instance assumes control of all relevant VIP addresses being used at that time by the PowerVille LB.

The following diagram shows the PowerVille LB utilizing the VIP Manager.



In the following diagram, the Active VIP Manager has failed. The Passive VIP Manager from the previous diagram takes over the Active VIP Manager role.

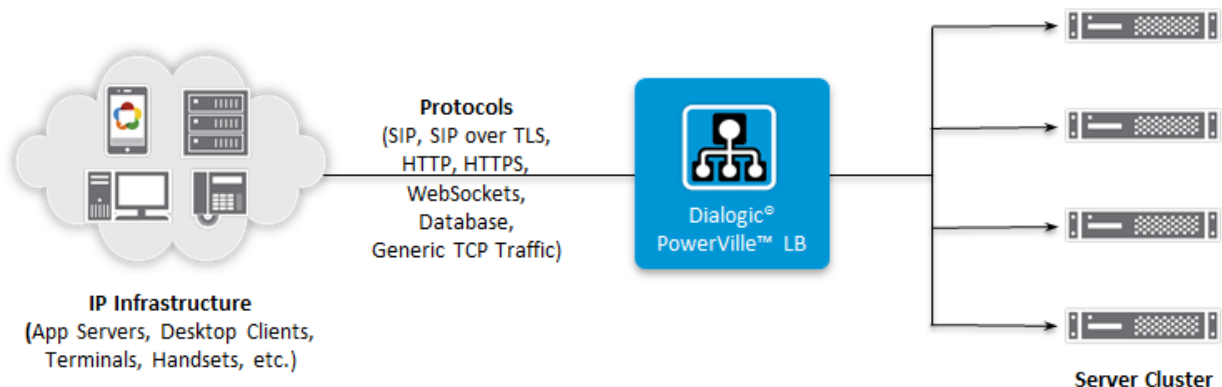


PowerVille LB Console

The PowerVille LB console is a web-based graphical user interface (WebGUI) used to manage the PowerVille LB. The PowerVille LB uses the Java Management Extensions (JMX) interface. The JMX interface is exposed by the Bootstrap Manager process and is accessed by a web application that runs in a web container. The Jetty web container is used to host the PowerVille LB web application. For more information on using the PowerVille LB console, refer to the *Dialogic® PowerVille™ LB – Load Balancer for Real-Time Communications Installation and Operations Guide*.

5. PowerVille LB Functionality

The PowerVille LB solution has many elements that allow it to meet network requirements including VIP addresses, selection algorithms, and HA, which are discussed in this section. The following diagram provides an overview of the PowerVille LB distributing traffic to a cluster of back-end servers.



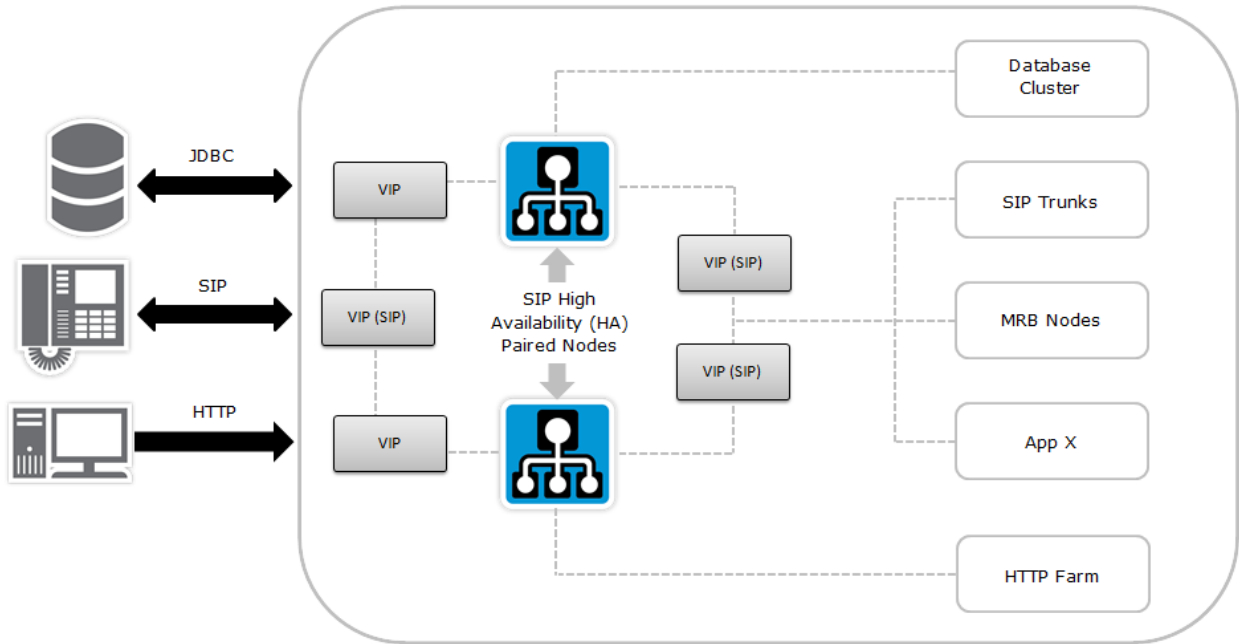
The left side of the diagram represents the ingress point presented by the PowerVille LB to solution clients. The clients use the VIP offered by the PowerVille LB to direct incoming traffic. The supported protocols include SIP, SIP over TLS, HTTP, HTTPS, WebSockets, Database, and generic TCP traffic. Each protocol type presents a different network port when offering service to the network. For example, on a single VIP address SIP can utilize port 5060, HTTP can utilize port 80, etc. Client traffic is then routed to an appropriate back-end server as represented by the right side of the diagram. The PowerVille LB determines the most appropriate back-end server based on the server's availability and the selection algorithm chosen for the service.

The PowerVille LB supports the following selection algorithms:

- **Round Robin** - The PowerVille LB loops through the available back-end servers in turn to choose the most appropriate one.
- **Priority** - The PowerVille LB uses an ordered list. The first available back-end server on the list is always used.
- **Prioritized Round Robin** - This is a combination of both the Priority and Round Robin algorithms. When this algorithm is enabled, the first entry in the services list of nodes takes priority and all traffic will be directed to this node. Should the node fail or be taken out of service, then the remaining entries in the list will be used in a Round Robin fashion.

After the PowerVille LB determines the most appropriate back-end server, traffic is dispatched to the selected server for processing.

The following diagram illustrates how the PowerVille LB services RTC using HA and VIPs to support clients.



In the diagram, there are two instances of the PowerVille LB that have been paired to provide an HA solution. Deploying in this manner is well-suited for RTC deployments. HA provides service continuity in the event that the primary PowerVille LB (the PowerVille LB node consuming traffic) suffers a catastrophic failure. To configure HA, two independent instances of PowerVille LB are required. The two instances are configured as paired nodes using the PowerVille LB console, which is controlled by a designated PowerVille LB instance. Once HA has been enabled and the PowerVille LBs have been provisioned, configuration information is automatically shared between the pair of PowerVille LB instances to provide a cluster. The PowerVille LB pairs interact with the VIP Manager to present a floating IP address. The floating IP address, or Virtual IP (VIP) address, is provisioned using the PowerVille LB console. In collaboration with a specified network port, the VIP is presented to the network to front the PowerVille LB pair. If the PowerVille LB node receiving traffic fails, the alternative PowerVille LB instance takes control of the VIP and continues service.

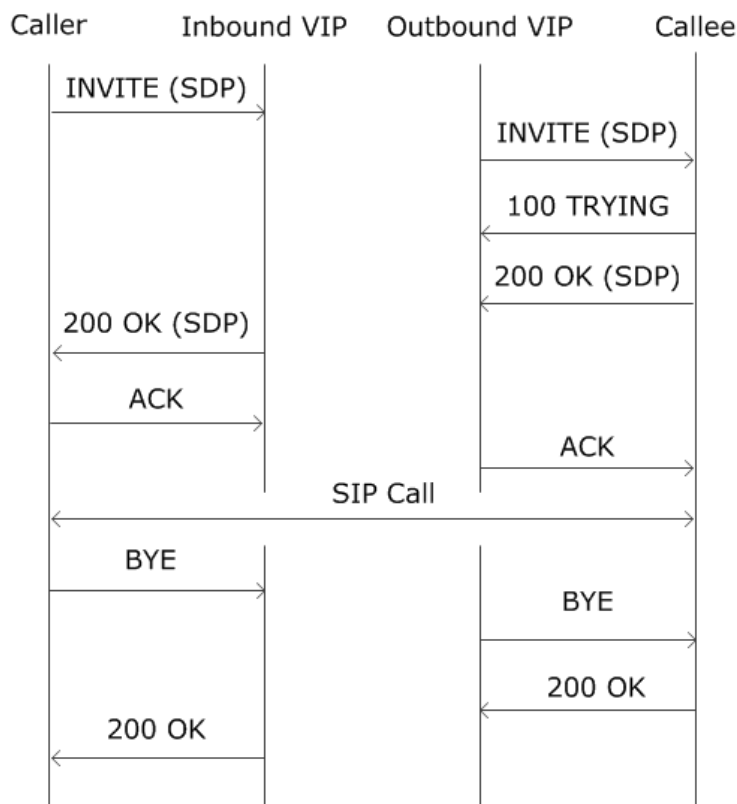
In addition to HA, the diagram also illustrates how SIP clients and SIP nodes interact when utilizing the PowerVille LB. The SIP protocol is inherently bidirectional so the PowerVille LB must fully support traffic initiated by both clients and server nodes. For example, on the left side of the diagram is a SIP client, which sends inbound messages to the PowerVille LB's "inbound" VIP. On the right side of the diagram are SIP nodes, which must be able to send outbound messages back to the SIP client via the "outbound" VIP. As such, the PowerVille LB supports inbound and outbound VIPs to facilitate the bidirectional transfer of SIP messaging. Unlike SIP, HTTP is not bidirectional, which is why only a single VIP is required for these services.

6. Message Flows

The PowerVile LB supports a number of complex deployment models and protocol level interactions. An important concept to keep in mind when analyzing the PowerVile LB is the usage of VIP addresses. For any type of protocol service exposed to the network, a VIP address is provisioned along with a unique service port. The VIP, in conjunction with the service port, fulfills the role of ingress point to the service for all incoming traffic. For example, a remote SIP client would be provisioned to send all its traffic to the provisioned VIP address and service port. This ingress point is known as the inbound VIP. The SIP protocol, however, adds some complexity because it is a bidirectional protocol that supports messaging in both directions (unlike HTTP). For a SIP service, an additional VIP is required to provide an egress point for messages traversing the PowerVile LB in the reverse direction (initiated by the back-end service node). This egress point is known as the outbound VIP. Refer to the [VIP Manager Process and High Availability](#) section for more information on VIP management.

SIP Call Flow

When a basic SIP call is directed to a service, the messages are sent to the inbound VIP address of the PowerVile LB. The PowerVile LB then selects an appropriate back-end service node depending on configuration and the selection algorithm provisioned. The PowerVile LB directs the call via the outbound VIP address to the receiving back-end service node. Refer to the following call flow diagram for details.



HTTP Message Flow

When considering a basic HTTP message flow example, the HTTP client makes an HTTP GET request to the inbound VIP address. The PowerVile LB then selects an appropriate back-end service node depending on configuration and the selection algorithm provisioned. The PowerVile LB directs the call via the PowerVile LB IP address to the receiving back-end service node. Since the HTTP protocol is not bidirectional like SIP, the PowerVile LB dispatches the HTTP GET request using its local PowerVile LB IP address instead of the outbound VIP address as done in the previous SIP example. Refer to the following call flow diagram for details.

