ORACLE
COMMUNICATIONS

SLB SBC Configuration Guidelines

Revision History

Version	Author	Description of Changes	Date Revision Completed
520-0045-00	Patrick Timmons	Initial Release	01/20/2012
520-0045-01	Soumil Vora	Update formatting and show running configs	08/01/2013
520-0045-02	Bhaskar Reddy Gaddam	Rebranding with latest release information	07/12/2018

Status of this memo

Oracle Corporation Best Current Practices are working documents of the Professional Services department of Oracle Corporation, Inc. Note that other groups may also distribute working documents as Best Current Practices.

Best Current Practices are working documents valid until explicitly obsoleted, and may be updated, replaced or obsoleted by other documents at any time. It is recommended to use Best Current Practices as reference material as well as to cite them in other works in progress.

Abstract

The use of the RFC 2119 keywords is an attempt to assign the correct requirement levels ("MUST", "SHOULD", "MAY", etc.).

This document defines a series of recommendations for Session-Aware Load Balancer (SLB) and SBC configuration on the Oracle Corporation AP6100 in a customer's production network. They should be used when either (a) deploying a new SLB/SBC, or (b) updating an existing configuration made before Best Current Practices were in place. When in conflict with Customer requirements or desires, the Customer's preference SHOULD take precedence.

Applicability

This document is applicable to AP6100 (running 7.3.0 or newer) in the role of the Session-Aware Load Balancer (SLB) and NN6100/NN4600 (running S-CZ8.1.0 or newer) Session Directors.

Contents

1 Introduction

As service providers deploy larger and larger SIP access networks, scalability problems are presenting unique challenges, particularly from an operational standpoint. Deployments that scale beyond the number of users serviceable by a single Session Border Controller (SBC) – as well as deployments that use a geographically redundant SBC for catastrophic fail over purposes – encounter edge reachability problems. In general there are two coarse techniques that carriers use today to support end-point populations that exceed one SLB's capacity: they either use a DNS-based distribution mechanism, or they will pre-provision endpoint to point to specific SBCs (manually load balancing them). Each of these solutions has its drawbacks. End users – many of them familiar with load balancing equipment deployed to scale protocols such as HTTP or SMTP – have expressed interest in a device that will perform dedicated load balancing for their SIP endpoint.

The Subscriber-Aware Load Balancer (SLB) addresses the need for scaling a network edge to millions of endpoint. Designed as a standalone system, the network architect can deploy an Acme Packet 6100, capable of supporting up to ten million endpoints (where an endpoint is defined as a unique source and destination IP address), the SLB aggregates signaling from large endpoint populations to reduce the edge reachability problem by an order of magnitude.

The network architect reduces this problem by deploying clusters of SBCs or Oracle Communications Unified Session Managers (USM) supported by the SLB. These SBCs can be operating as either Physical Network Functions (PNFs) and Virtual Network Functions (VNFs). The SLB supports clusters of homogenous or heterogenous groups of PNFs and/or VNFs.
.

Configuration guides are available for download from the Oracle Corporation Customer Support Portal (https://docs.oracle.com). Please contact your Oracle Corporation Systems Engineer for other Best Current Practice (BCP) documentation.

Functional Overview:

The Subscriber-Aware Load Balancer (SLB) is a discrete network element that processes all SIP end-point signaling traffic entering the service provider network. The SLB is not necessarily the first network device to receive signaling traffic, as, depending on network topology, additional network components (for example, routers, network address translators, and so on) can lie between the end-point and the SLB.
Upon receipt of a SIP packet from an unknown source, the SLB uses a provisioned policy to select an appropriate next-hop Session Border Controller (SBC) for traffic originated by that end-point. Subsequent packets from the same end-point are forwarded to the same SBC. The first packet, the one used to make the route decision, and all subsequent packets sent through the SLB to the next-hop SBC are encapsulated within an IP-in-IP format as defined in RFC 2003, IP Encapsulation within IP.
SBCs that participate in the load balancing-enabled deployment are enhanced by several capabilities. First, the SBC supports RFC 2003 tunnel for both packet transmission and reception. Second, the SBC periodically transmits health and performance data to the SLB; such information is evaluated and entered into the SLB's route determination algorithm. Lastly, the SBC participates in any SLB-initiated rebalance operation, as described in the Rebalancing section. A group of SBCs, with the above-listed capabilities, that receive signaling traffic from the SLB, is referred to as a cluster.

2 Intended Audience

This document is intended for use by Oracle Corporation Systems Engineers, third party Systems Integrators, and end users of the Session-Aware Load Balancer and Session Director. It assumes that the reader is familiar with basic operations of Acme Packet's ACLI, and has attended the following training courses (or has equivalent experience):

https://docs.oracle.com/cd/E95619_01/html/esbc_ecz810_configuration/

It also presumes that the reader is familiar with standard configuration models and archetypes; for more information, published in our Best Current Practice series of documentation.

3 Background

The popularity of consumer VoIP services has stimulated a rapid growth of carriers' networks to support those services. As the number of devices attaching to a network climbs, the "edge discovery" problem for those devices climbs commensurately. Session Border Controllers, such as Acme Packet's Session Director, are logically situated at the border between a public network (e.g., the Internet) and a private carrier network. These SBCs have a finite capacity for handling user traffic, and therefore each member of the growing user population needs to select from one of many possible border points into their service provider's network.

This raises some interesting design challenges for service providers; first, how do endpoints choose their SBC. Historically, the predominant technology for attaching endpoints to SBCs is to use DNS. While it is a venerable protocol (and therefore well understood, easy to manage, scalable, and presumably part of large IP networks already), it is not particularly well suited to SIP networks for a number of reasons. Endpoint implementation decisions aside, using DNS to make distribution decisions is non-deterministic; that is, any number of intermediaries between a DNS client and a DNS server can cache data, thereby losing granular control over when services advertised or withdrawn by the DNS server are effective. Further, while DNS can use information about the requestor to influence its response, most DNS distribution algorithms are extremely rudimentary (round-robin).

The second design challenge carriers face is how to manage the user attachment to the network, for troubleshooting or sizing purposes. When investigating a customer issue, it is impractical to hunt through all potential entry points for state data related to a specific user. And as the capacity on existing systems rises, adding a system to increase overall service capacity, while redistributing users among the increased pool of resources without service disruption is a delicate, costly operation.

The Oracle Corporation clustering model using the SLB is designed to scale access networks up to two million individual connections, while eliminating dependencies on the endpoint (DNS behavior, 3xx redirect behavior, etc.) and providing a clean, simple, adaptable interface into and out of a carrier's network. By using an SLB as the entry point into a cluster of back-end SBC equipment, the edge discovery problem is reduced by an order of magnitude, and the SLB's flexible policy-driven distribution architecture puts endpoints where they belong – even if the network topology changes. The SLB and SDs form a cohesive unit, relaying performance and platform data to optimize the distribution of resources among all devices.

4 Design Goals

The intents of the configurations in this document are to:

- Minimize interoperability issues by standardizing field configurations
- Provide guidelines for new users to the Session-Aware Load Balancer and Session Director
- Document when and why configuration elements should be changed from their default values
- Facilitate transition of customers from Systems Engineering to Technical Support by making configurations consistent (yielding predictable behavior)
- Illustrate how to transition existing Session Directors into a cluster fronted by a Session-Aware Load Balancer by highlighting the configuration changes required from standard, non-clustered BCP examples.

Further, each design considers the following aspects (in order of priority):

- Flexibility: how resilient the configuration is, and how adaptable the configuration is when turning up new connected networks (for example)
- Performance: minimizing the use of “heavy” configuration objects, to streamline the message flow through the box and reduce CPU usage.
- Scalability: minimizing redundant configuration objects and setting a template based foundation to allow overlay configuration with minimal disruption
- Compatibility: working with other popular devices in carriers' VoIP networks

5 Notes on the Reference Configurations

All of the configurations presented here have been entered, tested, and verified on a SD in the lab at headquarters. The goal is not to demonstrate a full-featured configuration; rather, each contains only the minimum number of configuration objects required to pass basic SIP transactions. The appendices of this document provide a single working, tested instance of each of the SLB and supporting SD configurations. Only the default SLB distribution policy (i.e., round-robin) is employed by this sample.

In all cases, the design uses a single "untrusted" network, in the 192.168.11.0/24 subnet, and a single "trusted" network using 192.168.12.0/24. The configurations have been designed such that no 192.168.12.0/24 IP addresses are leaked into signaling messages sent to the untrusted network. Additionally, a "tunnel" network exists in cluster configurations that are not present in standard SBC applications. This is a network that exists in the signaling plane between an SLB and the SBCs in its cluster. The sample tunnel network in this document uses a 172.16.0.0/16 private network space, although please note this tunnel network may exist over any routable IP network (public or private, localized or geographically spread).

The IP address to which UAs send their SIP signaling in all cases is 192.168.11.101 on the SLB; this is represented as (a) in the figure below. Note that this address is also present in the sample SD configuration and represented in the diagram as address (c); however, the SD does not advertise this address on its network. The SD only uses this address via the IP-in-IP tunnel established between it and its SLB. The tunnel IP addresses, (b) and (d) in the diagram, are on the tunnel network, and are 172.16.0.100 and 172.16.0.200, respectively. The IP address (e) from which the SD sends its messages to the core infrastructure is 192.168.12.100. Depending on the configuration model, there may be more addresses used on either network; this will be noted as applicable. The SIP registrar used for testing is located at 192.168.12.200, depicted as address (f).

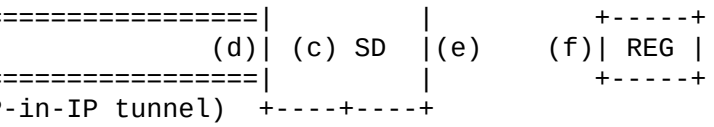

No Denial of Service (DoS) configuration has been applied on the SD's configuration, save the application-layer access control features of the sip-interface. Only rudimentary DoS configuration is applied to the SLB, in the form of its max-untrusted-percentage. (This value represents the maximum percentage of the SLB's forwarding table that is available to untrusted/temporary entries.) NOTE: in its present state, the interface used for CCP messaging is instantiated as an "untrusted" flow; this means that it is possible to overwhelm a cluster with a DoS attack, as congestion on the interface between the SLB and SD will cause keep-alive signaling to be throttled – causing the SLB to believe that the SD has failed. It is therefore STRONGLY RECOMMENDED to configure a static, trusted ACL on the SLB for each SD in the cluster (or a subnetted ACL that covers them, etc.).

Note that registration-caching must be ENABLED on the appropriate sip interface when clustering an SD. The SD uses this cached registration data for organizing and prioritizing its endpoint data during cluster rebalancing operations.

The systems used for testing purposes were configured in a high availability environment. For more information on configuring High Availability on the Session Director, refer to the pertinent BCP documentation for the latest best practices. For more information on configuring High Availability on the Session-Aware Load Balancer, refer to Section 6 of this document.

Best Current Practices for object naming conventions have been followed whenever possible.

6 Session-Aware Load Balancer Configuration

The SLB configuration is deliberately kept simple. Aside from the usual array of "platform" configuration elements (e.g., system-config, phy-config, etc.) one must only configure two global configuration settings, and any desired, specific distribution policies. (Note that in the absence of any configured distribution policies, the SLB will default to a round-robin mode of operation, which is what the sample configuration presents.)

6.1. The cluster-config element

The default values of the cluster-config element have been chosen to be applicable to the majority of deployment scenarios. Various configuration settings should be specifically considered when deploying the SLB, however.

The boolean value "auto-rebalance" governs whether or not the addition of a new SD to a cluster will trigger a rebalance operation within that cluster (i.e., each pre-existing SD will migrate some percentage of their user population to fill this new vacancy). Should a carrier's preference be to keep traffic static upon node insertion, this value should be set to disabled; the result of this configuration will be that the newly added SD is considered the de facto highest preference for new users (so it should fill up rapidly). In most cases auto-rebalance should be set to enabled.

The rebalance-skip-ahead value (represented in milliseconds) controls how far into its time-ordered list of endpoint expiry times the SD will skip to find candidates for rebalancing to other SDs when asked to nominate candidates by the SLB. By default (and when the rebalance-skip-ahead value is set to 0) the SD will nominate users at the "top" of its time-ordered list; that is, the users expiring most imminently will be the ones chosen. However, this introduces a distinct possibility for a race condition, where the SD nominates a user whose refresh REGISTER is already 'in flight' towards the SLB. By setting this field to a non-zero value (such as the recommended value of thirty seconds, or 30000), the SD will nominate users set to expire thirty seconds from the time of the request from the SLB. This significantly mitigates the likelihood of an SD nominating an endpoint about to send its REGISTER, introducing a state abnormality among the cluster members. The value of thirty seconds was selected deliberately to roughly coincide with the duration of a typical SIP REGISTER transaction. Likewise, the rebalance-max-refresh will allow an operator to establish a ceiling on the time differential between when an endpoint may be moved (deleted on the SLB) and when it is expected to refresh its registration through the SLB. As a move operation from one SD to another SD will necessarily introduce a window of unavailability for that endpoint to receive calls successfully via the cluster, it is oftentimes desirable to control that window explicitly. A carrier willing to expose their users to no more than five minutes of downtime during a cluster rebalance operation, for example, could set rebalance-max-refresh to 300000. In this way, only users due to expire five minutes from the receipt of the rebalance request from the SLB would be considered candidates for migration.

The session-multiplier field is used to derive the occupancy for a given Session Director in the duster. Occupancy is essentially the governing factor for which SD is chosen for a new endpoint assignment, when all other factors are equal. Said another way, if two systems are identically configured and a new endpoint signals into the cluster, the one with the lowest occupancy is the one that will be selected. When an SD first joins the cluster (handshakes into the SLB), it will provide several aspects of its configuration – among them its licensed session capacity, and its maximum endpoint capacity (if configured). If its maximum endpoint capacity is configured (sip-config -> registration-cache-limit), then this is considered the maximum occupancy of the SD. If it is not configured, the SLB uses the product of the licensed capacity and the session-multiplier field from the cluster-config. As an example, using the default session-multiplier of 10, an SD with 16,000 licensed sessions will receive no more than 160,000 endpoint assignments from the SLB. The default session-multiplier should be adequate for most applications, but it is important to evaluate it independently for each cluster deployment.

6.2. The lbp-config element

Generally speaking, no configuration changes (aside from any relative to high availability) are usually necessary in the lbp-config element; the default values are adequate.

However, certain deployment scenarios may justify changing the untrusted-grace-period timer from its default of 30 (seconds) to a higher value. In a typical access configuration, the untrusted sip-interface is configured to restrict its use to registered endpoints only (allow-anonymous->registered). There are some instances, however, where this security feature is either not applied, or is deliberately overridden by configuration specific to emergency calling.

In its present incarnation, the SLB will only migrate a user from its 'untrusted' list to its 'trusted' list via a promotion message sent by the SD specific to that user; in turn, today's SBC incarnation will only send this promotion message to the SLB after a successful REGISTER/200 OK exchange has occurred. The net result of this is that an anonymous caller, sending an INVITE to an SD where the user is not currently registered, will complete their call but the SLB will never 'trust' that user. After the untrusted-grace-period elapses, the SLB will 'forget' about that particular endpoint-to-SD association, and the user's subsequent signaling messages will be subject to new distribution decisions. Adjusting the untrusted-grace-period timer value to match the user's registration interval will allow the association to exist long enough for that anonymous caller to re-REGISTER during their call, and hence get promoted by the SD after the REGISTER/200 OK. (Note: even if a user gets redistributed to a different SD during an active call it should not interrupt their call in progress, as RTP never flows through the SLB; the redistribution would only affect in-dialog signaling messages such as the BYE, etc.)

6.3. Configuring distribution policies (lb-policy)

The SLB's distribution policies, configured using lb-policy configuration objects, are where customer-specific mappings from endpoints to supporting SDs are defined. As such, no specific configuration advice is pertinent in this BCP.

In the absence of any configured lb-policy objects, the SLB's default behavior is to round-robin all inbound requests among all SDs that advertise support for the service interface upon which the request arrived at the SLB. If round-robin behavior is all that is desired, no lb-policy elements SHOULD be configured. However, any advanced routing will require the use of one or more lb-policy configuration objects to define the desired mapping of inbound endpoints to SDs or groups of SDs.

Fundamentally, lb-policy configuration is roughly analogous to local-policy on the SD; that is, it will match packets to policies using their source address and destination address, prioritize the results, and forward the packet onto the "best" match. Unlike local-policy matching, however, the "source realm" is not part of the matching key on ingress packets, since the SLB is neither aware of, nor configured with, realm data. Combining the local-policy concepts of a 'next-hop' and an (egress) realm within a policy-attribute is the name of a "service partition", defined as the realm identifier on each member SD. Each of these service partitions is configured as an "lb-realm", a sub-element within the lb-policy configuration element. The name of the lb-realm matches the name of one or more realm labels (realm-config -> identifier) of participating SDs.

6.4. High availability configuration

The SLB is responsible for replicating data pertaining to as many as two million endpoints. As such, those familiar with configuring high availability on an SD will notice some recommendations for changing the default values of various timers and transaction journal settings.

In the lbp-config, red-max-trans governs the journal depth for queuing transaction history on the active SLB device. That is, each transaction that occurs on the active SLB is recorded in this journal awaiting a request from the standby SLB to retrieve updates. For scaling up to two million supported endpoints, the recommended value for this configuration setting is 2000000. This will allow for a scenario where the high availability mate of an active SLB is powered off for an extended period of time; when the high availability pair is restored, the journal on the active will be large enough to accommodate a (potentially) large number of transactions that have transpired in the interim.

In the redundancy-config, the recommended default value for becoming-standby-time is 1620000 (27 minutes). This is to account for the large volume of data a standby can expect to receive from an active, and gives it time to retrieve all of that data and become fully synchronized before this timer expires and it transitions to OutOfService.

The sample configuration is given in Appendix A.

7 Session Director Configuration

The configuration requirements for introducing a Session Director into a cluster fronted by an SLB were deliberately kept minimal. All that is required is to configure the properties of the tunnel between the SD and SLB, and to flag the relevant application(s) to indicate their participation in a cluster.

7.1. Tunnel configuration

Traditionally, SDs are situated at the border of two networks: an untrusted network and a trusted network. The introduction of an SLB into a network design creates a third network, referred to as a "tunnel network". (Therefore the SLB is said to participate in the untrusted network and the tunnel network, whereas the SDs in the cluster participate in all three of the untrusted, tunnel, and trusted networks.) As the tunnel properties are defined with a network-interface, any realms built upon that network-interface that are configured to communicate with a load balancer (see next section) will use the tunnel. Multiple tunnels may be configured on a single network-interface (differentiated by their name), but only one tunnel per protocol type (e.g., SIP, H248) may be configured within a single network-interface.

7.2. Application configuration

As the application configuration consists solely of either a tunnel identifier field in the sip-interface, or an option within the h248-mgc-config element, no additional information is necessary.

Note that when configuring an SD to support SIP/TCP or SIP/TLS, that the option "reuse-connections" MUST be applied to the access-facing (i.e., tunneled) sip-interface. This is because an SD may try to initiate outbound connections to non-NATted endpoints if the IP:port in the endpoint's Contact-URI does not match its actual L3/L4 address. These connections will fail in the presence of a SLB.

7.3. Media management in a cluster

Be advised that media management (specifically, mm-in-realm) will behave differently as endpoints are strewn among many like-configured SDs. If a customer is intent on releasing media between two users within a realm, you must consider that the realm now extends to the entire cluster; therefore, you must configure msm-release in addition to setting mm-in-realm to disabled.

For more information on configuring msm-release, refer to [2].

8 Maintenance and Troubleshooting

Please refer to the document available at the link below from our support portal for helpful show command outputs and their explanations:-

https://docs.oracle.com/cd/E83022_01/doc/slb_scz7310_essentials.pdf

9 References

1. https://docs.oracle.com/cd/E83022_01/index.htm

10 Author´s Address

Bhaskar Reddy Gaddam
Oracle Communications.
100 Crosby Drive
Bedford, MA 01730
Email: bhaskar.gaddam@oracle.com

11 Disclaimer

The content in this document is for informational purposes only and is subject to change by Oracle Corporation without notice. While reasonable efforts have been made in the preparation of this publication to assure its accuracy, Oracle Corporation assumes no liability resulting from technical or editorial errors or omissions, or for any damages resulting from the use of this information. Unless specifically included in a written agreement with Oracle Corporation, Oracle Corporation has no obligation to develop or deliver any future release or upgrade or any feature, enhancement or function.

12 Full Copyright Statement

APPENDIX A. Reference Configuration: SLB Basic Configuration

```
cluster-config
        state                           enabled
        log-level                       CRITICAL
        auto-rebalance                  enabled
        source-rebalance-threshold      50
        dest-rebalance-threshold        0
        dest-rebalance-max              80
        tunnel-check-interval           15000
        tunnel-fail-interval            10000
        rebalance-request-delay         500
        session-multiplier              10
        rebalance-skip-ahead            30000
        rebalance-max-refresh           0
        ignore-tgt-svcs-on-rebalance    disabled
        atom-limit-divisor              1
        rebalance-del-app-entries       disabled
        inactive-sd-limit               1800
        red-port                        2001
        red-max-trans                   10000
        red-sync-start-time             5000
        red-sync-comp-time              1000
        service-port
                address                         192.168.11.101
                port                            5060
                protocol                        UDP
                network-interface               M00:0
        last-modified-by                admin@console
        last-modified-date              2013-07-03 14:12:44
ipt-config
        min-tru-bw                      8000000
        max-tru-bw                      24000000
        min-untru-bw                    3000000
        max-untru-bw                    9000000
        log-level                       CRITICAL
        last-modified-by                admin@console
        last-modified-date              2013-07-03 14:19:45
lbp-config
        state                           enabled
        log-level                       CRITICAL
        untrusted-grace-period          30
        max-untrusted-percentage        20
        max-untrusted-upper-threshold   80
        max-untrusted-lower-threshold   70
        endpoint-capacity-upper-threshold 80
        endpoint-capacity-lower-threshold 70
        red-port                        2000
        red-max-trans                   2000000
        red-sync-start-time             5000
        red-sync-comp-time              1000
        last-modified-by                admin@console
        last-modified-date              2013-07-03 14:13:53
network-interface
        name                            M00
        sub-port-id                     0
        description                   Access
        hostname
        ip-address                      192.168.11.101
        pri-utility-addr                192.168.11.102
        sec-utility-addr                192.168.11.103
```



```
        netmask                         255.255.255.0
        gateway                         192.168.11.1
        sec-gateway
        gw-heartbeat
                state                           disabled
                heartbeat                       0
                retry-count                     0
                retry-timeout                   1
                health-score                    0
        dns-ip-primary
        dns-ip-backup1
        dns-ip-backup2
        dns-domain
        dns-timeout                     11
        hip-ip-list
        ftp-address
        icmp-address
        snmp-address
        telnet-address
        ssh-address
        last-modified-by                admin@console
        last-modified-date              2013-07-03 14:05:40
network-interface
        name                            M10
        sub-port-id                     0
        description                     tunnel
        hostname
        ip-address                      172.16.0.100
        pri-utility-addr                172.16.0.101
        sec-utility-addr                172.16.0.102
        netmask                         255.255.255.0
        gateway                         172.16.0.1
        sec-gateway
        gw-heartbeat
                state                           disabled
                heartbeat                       0
                retry-count                     0
                retry-timeout                   1
                health-score                    0
        dns-ip-primary
        dns-ip-backup1
        dns-ip-backup2
        dns-domain
        dns-timeout                     11
        hip-ip-list                     172.16.0.100
        ftp-address
        icmp-address                    172.16.0.100
        snmp-address
        telnet-address
        ssh-address
        tunnel-config
                local-address                   172.16.0.100
                port                            4444
                protocol                        UDP
                tls-profile
        last-modified-by                admin@console
        last-modified-date              2013-07-03 14:07:35
network-interface
        name                            eth1
        sub-port-id                     0
        description                     wancom1
        hostname
        ip-address
        pri-utility-addr                169.254.1.1
```



```
        sec-utility-addr                169.254.1.2
        netmask                         255.255.255.252
        gateway
        sec-gateway
        gw-heartbeat
                state                           disabled
                heartbeat                       0
                retry-count                     0
                retry-timeout                   1
                health-score                    0
        dns-ip-primary
        dns-ip-backup1
        dns-ip-backup2
        dns-domain
        dns-timeout                     11
        hip-ip-list
        ftp-address
        icmp-address
        snmp-address
        telnet-address
        ssh-address
        last-modified-by                admin@console
        last-modified-date              2013-05-08 11:38:14
network-interface
        name                            eth2
        sub-port-id                     0
        description                     wancom2
        hostname
        ip-address
        pri-utility-addr                169.254.2.1
        sec-utility-addr                169.254.2.2
        netmask                         255.255.255.252
        gateway
        sec-gateway
        gw-heartbeat
                state                           disabled
                heartbeat                       0
                retry-count                     0
                retry-timeout                   1
                health-score                    0
        dns-ip-primary
        dns-ip-backup1
        dns-ip-backup2
        dns-domain
        dns-timeout                     11
        hip-ip-list
        ftp-address
        icmp-address
        snmp-address
        telnet-address
        ssh-address
        last-modified-by                admin@console
        last-modified-date              2013-05-08 11:38:54
phy-interface
        name                            M00
        operation-type                  Media
        port                            0
        slot                            0
        virtual-mac                     00:08:25:a2:39:fe
        admin-state                     enabled
        auto-negotiation                enabled
        duplex-mode                     FULL
        speed                           100
        overload-protection             disabled
```



```
        last-modified-by                admin@console
        last-modified-date              2013-05-23 15:33:09
phy-interface
        name                            M10
        operation-type                  Media
        port                            0
        slot                            1
        virtual-mac                     00:08:25:a2:39:ff
        admin-state                     enabled
        auto-negotiation                enabled
        duplex-mode                     FULL
        speed                           100
        overload-protection             disabled
        last-modified-by                admin@console
        last-modified-date              2013-05-23 17:15:25
phy-interface
        name                            eth1
        operation-type                  Control
        port                            1
        slot                            0
        virtual-mac
        wancom-health-score             8
        overload-protection             disabled
        last-modified-by                admin@console
        last-modified-date              2013-05-08 11:44:37
phy-interface
        name                            eth2
        operation-type                  Control
        port                            2
        slot                            0
        virtual-mac
        wancom-health-score             9
        overload-protection             disabled
        last-modified-by                admin@console
        last-modified-date              2013-05-08 11:45:02
redundancy-config
        state                           enabled
        log-level                       INFO
        health-threshold                75
        emergency-threshold             50
        port                            9090
        advertisement-time              500
        percent-drift                   210
        initial-time                    1250
        becoming-standby-time           180000
        becoming-active-time            100
        cfg-port                        1987
        cfg-max-trans                   10000
        cfg-sync-start-time             5000
        cfg-sync-comp-time              1000
        gateway-heartbeat-interval      0
        gateway-heartbeat-retry         0
        gateway-heartbeat-timeout       1
        gateway-heartbeat-health        0
        media-if-peercheck-time         0
        peer
                name                            SLB1
                state                           enabled
                type                            Primary
                destination
                        address                         169.254.1.1:9090
                        network-interface               eth1:0
                destination
                        address                         169.254.2.1:9090
```



```
                        network-interface               eth2:0
        peer
                name                            SLB2
                state                           enabled
                type                            Secondary
                destination
                        address                         169.254.2.2:9090
                        network-interface               eth2:0
                destination
                        address                         169.254.1.2:9090
                        network-interface               eth1:0
        last-modified-by                admin@console
        last-modified-date              2013-05-08 13:54:12
phy-interface
        name                            M10
        operation-type                  Media
        port                            0
        slot                            1
        virtual-mac
        admin-state                     enabled
        auto-negotiation                enabled
        duplex-mode                     FULL
        speed                           100
        overload-protection             disabled
        last-modified-by                admin@console
        last-modified-date              2013-07-03 13:57:47
phy-interface
        name                            M00
        operation-type                  Media
        port                            0
        slot                            0
        virtual-mac
        admin-state                     enabled
        auto-negotiation                enabled
        duplex-mode                     FULL
        speed                           100
        overload-protection             disabled
        last-modified-by                admin@console
        last-modified-date              2013-07-03 13:58:01
system-config
        hostname
        description                     SLB BCP Config - LCX 1.0.0
        location
        mib-system-contact
        mib-system-name
        mib-system-location
        snmp-enabled                    enabled
        enable-snmp-auth-traps          disabled
        enable-snmp-syslog-notify       disabled
        enable-snmp-monitor-traps       disabled
        enable-env-monitor-traps        disabled
        snmp-syslog-his-table-length    1
        snmp-syslog-level               WARNING
        system-log-level                WARNING
        process-log-level               NOTICE
        process-log-ip-address          0.0.0.0
        process-log-port                0
        collect
                sample-interval                 5
                push-interval                   15
                boot-state                      disabled
                start-time                      now
                end-time                        never
                red-collect-state               disabled
```



```
                red-max-trans                   1000
                red-sync-start-time             5000
                red-sync-comp-time              1000
                push-success-trap-state         disabled
        call-trace                      disabled
        internal-trace                  disabled
        log-filter                      all
        default-gateway                 172.41.0.1
        restart                         enabled
        exceptions
        telnet-timeout                  0
        console-timeout                 0
        remote-control                  enabled
        cli-audit-trail                 enabled
        link-redundancy-state           disabled
        source-routing                  disabled
        cli-more                        disabled
        terminal-height                 24
        debug-timeout                   0
        trap-event-lifetime             0
        default-v6-gateway              ::
        ipv6-support                    disabled
        cleanup-time-of-day             00:00
        last-modified-by                admin@console
        last-modified-date              2013-07-03 13:56:33
task done
```


APPENDIX B. Reference Configuration: Basic SD Clustered Configuration

```
local-policy
        from-address
                                        *
        to-address
                                        *
        source-realm
                                        Access
        description
        activate-time                   N/A
        deactivate-time                 N/A
        state                           enabled
        policy-priority                 none
        last-modified-by                admin@172.41.1.2
        last-modified-date              2013-07-03 11:15:59
        policy-attribute
                next-hop                        192.168.12.200
                realm                           Core
                action                          none
                terminate-recursion             disabled
                carrier
                start-time                      0000
                end-time                        2400
                days-of-week                    U-S
                cost                            0
                app-protocol
                state                           enabled
                methods
                media-profiles
                lookup                          single
                next-key
                eloc-str-lkup                   disabled
                eloc-str-match
media-manager
        state                           enabled
        latching                        enabled
        flow-time-limit                 86400
        initial-guard-timer             300
        subsq-guard-timer               300
        tcp-flow-time-limit             86400
        tcp-initial-guard-timer         300
        tcp-subsq-guard-timer           300
        tcp-number-of-ports-per-flow    2
        hnt-rtcp                        disabled
        algd-log-level                  NOTICE
        mbcd-log-level                  NOTICE
        red-flow-port                   1985
        red-mgcp-port                   1986
        red-max-trans                   10000
        red-sync-start-time             5000
        red-sync-comp-time              1000
        media-policing                  enabled
        max-signaling-bandwidth         10000000
        max-untrusted-signaling         100
        min-untrusted-signaling         30
        app-signaling-bandwidth         0
        tolerance-window                30
        rtcp-rate-limit                 0
        trap-on-demote-to-deny          disabled
        anonymous-sdp                   disabled
```



```
        arp-msg-bandwidth               32000
        fragment-msg-bandwidth          0
        rfc2833-timestamp               disabled
        default-2833-duration           100
        rfc2833-end-pkts-only-for-non-sig enabled
        translate-non-rfc2833-event     disabled
        media-supervision-traps         disabled
        dnsalg-server-failover          disabled
        last-modified-by                admin@172.41.1.2
        last-modified-date              2012-01-12 20:04:48
network-interface
        name                            M00
        sub-port-id                     0
        description                     Access Interface
        hostname
        ip-address                      172.16.0.200
        pri-utility-addr
        sec-utility-addr
        netmask                         255.255.255.0
        gateway                         172.16.0.1
        sec-gateway
        gw-heartbeat
                state                           disabled
                heartbeat                       0
                retry-count                     0
                retry-timeout                   1
                health-score                    0
        dns-ip-primary
        dns-ip-backup1
        dns-ip-backup2
        dns-domain
        dns-timeout                     11
        hip-ip-list                     172.16.0.200
        ftp-address
        icmp-address                    172.16.0.200
        snmp-address
        telnet-address
        ssh-address
        signaling-mtu                   0
        tunnel-config
                name                            sipTunnel
                local-address                   172.16.0.200
                remote-address                  172.16.0.100
                port                            4444
                protocol                        UDP
                tls-profile
                application                     SIP
        last-modified-by                admin@172.41.1.2
        last-modified-date              2013-07-03 14:29:08
network-interface
        name                            M10
        sub-port-id                     0
        description                     Core Interface
        hostname
        ip-address                      192.168.12.100
        pri-utility-addr
        sec-utility-addr
        netmask                         255.255.255.0
        gateway                         192.168.12.1
        sec-gateway
        gw-heartbeat
                state                           disabled
                heartbeat                       0
                retry-count                     0
```



```
                retry-timeout                  1
                health-score                   0
        dns-ip-primary
        dns-ip-backup1
        dns-ip-backup2
        dns-domain
        dns-timeout                    11
        hip-ip-list                    192.168.12.100
        ftp-address
        icmp-address                   192.168.12.100
        snmp-address
        telnet-address
        ssh-address
        signaling-mtu                  0
        last-modified-by               admin@console
        last-modified-date             2013-05-08 15:47:12
phy-interface
        name                           M00
        operation-type                 Media
        port                           0
        slot                           0
        virtual-mac
        admin-state                    enabled
        auto-negotiation               enabled
        duplex-mode                    FULL
        speed                          1000
        overload-protection            disabled
        last-modified-by               admin@172.41.1.2
        last-modified-date             2012-01-12 19:55:26
phy-interface
        name                           M10
        operation-type                 Media
        port                           0
        slot                           1
        virtual-mac
        admin-state                    enabled
        auto-negotiation               enabled
        duplex-mode                    FULL
        speed                          1000
        overload-protection            disabled
        last-modified-by               admin@172.41.1.2
        last-modified-date             2012-01-12 19:55:48
realm-config
        identifier                     Access
        description
        addr-prefix                    0.0.0.0
        network-interfaces
                                       M00:0
        mm-in-realm                    enabled
        mm-in-network                  enabled
        mm-same-ip                     enabled
        mm-in-system                   enabled
        bw-cac-non-mm                  disabled
        msm-release                    disabled
        qos-enable                     disabled
        generate-UDP-checksum          disabled
        max-bandwidth                  0
        fallback-bandwidth             0
        max-priority-bandwidth         0
        max-latency                    0
        max-jitter                     0
        max-packet-loss                0
        observ-window-size             0
        parent-realm
```



```
        dns-realm
        media-policy
        media-sec-policy
        srtp-msm-passthrough            disabled
        in-translationid
        out-translationid
        in-manipulationid
        out-manipulationid
        manipulation-string
        manipulation-pattern
        class-profile
        average-rate-limit              0
        access-control-trust-level      none
        invalid-signal-threshold        0
        maximum-signal-threshold        0
        untrusted-signal-threshold      0
        nat-trust-threshold             0
        deny-period                     30
        ext-policy-svr
        diam-e2-address-realm
        symmetric-latching              disabled
        pai-strip                       disabled
        trunk-context
        early-media-allow
        enforcement-profile
        additional-prefixes
        restricted-latching             none
        restriction-mask                32
        accounting-enable               enabled
        user-cac-mode                   none
        user-cac-bandwidth              0
        user-cac-sessions               0
        icmp-detect-multiplier          0
        icmp-advertisement-interval     0
        icmp-target-ip
        monthly-minutes                 0
        net-management-control          disabled
        delay-media-update              disabled
        refer-call-transfer             disabled
        refer-notify-provisional        none
        dyn-refer-term                  disabled
        codec-policy
        codec-manip-in-realm            disabled
        constraint-name
        call-recording-server-id
        xnq-state                       xnq-unknown
        hairpin-id                      0
        stun-enable                     disabled
        stun-server-ip                  0.0.0.0
        stun-server-port                3478
        stun-changed-ip                 0.0.0.0
        stun-changed-port               3479
        match-media-profiles
        qos-constraint
        sip-profile
        sip-isup-profile
        block-rtcp                      disabled
        hide-egress-media-update        disabled
        last-modified-by                admin@172.41.1.2
        last-modified-date              2012-01-12 20:05:08
realm-config
        identifier                      Core
        description
        addr-prefix                     0.0.0.0
```



```
        network-interfaces
                                      M10:0
        mm-in-realm                   disabled
        mm-in-network                 enabled
        mm-same-ip                    enabled
        mm-in-system                  enabled
        bw-cac-non-mm                 disabled
        msm-release                   disabled
        qos-enable                    disabled
        generate-UDP-checksum         disabled
        max-bandwidth                 0
        fallback-bandwidth            0
        max-priority-bandwidth        0
        max-latency                   0
        max-jitter                    0
        max-packet-loss               0
        observ-window-size            0
        parent-realm
        dns-realm
        media-policy
        media-sec-policy
        srtp-msm-passthrough          disabled
        in-translationid
        out-translationid
        in-manipulationid
        out-manipulationid
        manipulation-string
        manipulation-pattern
        class-profile
        average-rate-limit            0
        access-control-trust-level    none
        invalid-signal-threshold      0
        maximum-signal-threshold      0
        untrusted-signal-threshold    0
        nat-trust-threshold           0
        deny-period                   30
        ext-policy-svr
        diam-e2-address-realm
        symmetric-latching            disabled
        pai-strip                     disabled
        trunk-context
        early-media-allow
        enforcement-profile
        additional-prefixes
        restricted-latching           none
        restriction-mask              32
        accounting-enable             enabled
        user-cac-mode                 none
        user-cac-bandwidth            0
        user-cac-sessions             0
        icmp-detect-multiplier        0
        icmp-advertisement-interval   0
        icmp-target-ip
        monthly-minutes               0
        net-management-control        disabled
        delay-media-update            disabled
        refer-call-transfer           disabled
        refer-notify-provisional      none
        dyn-refer-term                disabled
        codec-policy
        codec-manip-in-realm          disabled
        constraint-name
        call-recording-server-id
        xnq-state                     xnq-unknown
```



```
        hairpin-id                      0
        stun-enable                     disabled
        stun-server-ip                  0.0.0.0
        stun-server-port                3478
        stun-changed-ip                 0.0.0.0
        stun-changed-port               3479
        match-media-profiles
        qos-constraint
        sip-profile
        sip-isup-profile
        block-rtcp                      disabled
        hide-egress-media-update        disabled
        last-modified-by                admin@console
        last-modified-date              2013-05-08 15:46:10
session-agent
        hostname                        192.168.12.200
        ip-address                      192.168.12.200
        port                            5060
        state                           enabled
        app-protocol                    SIP
        app-type
        transport-method                UDP
        realm-id                        Core
        egress-realm-id
        description
        carriers
        allow-next-hop-lp               enabled
        constraints                     disabled
        max-sessions                    0
        max-inbound-sessions            0
        max-outbound-sessions           0
        max-burst-rate                  0
        max-inbound-burst-rate          0
        max-outbound-burst-rate         0
        max-sustain-rate                0
        max-inbound-sustain-rate        0
        max-outbound-sustain-rate       0
        min-seizures                    5
        min-asr                         0
        time-to-resume                  0
        ttr-no-response                 0
        in-service-period               0
        burst-rate-window               0
        sustain-rate-window             0
        req-uri-carrier-mode            None
        proxy-mode
        redirect-action
        loose-routing                   enabled
        send-media-session              enabled
        response-map
        ping-method
        ping-interval                   0
        ping-send-mode                  keep-alive
        ping-all-addresses              disabled
        ping-in-service-response-codes
        out-service-response-codes
        load-balance-dns-query          hunt
        media-profiles
        in-translationid
        out-translationid
        trust-me                        disabled
        request-uri-headers
        stop-recurse
        local-response-map
```



```
        ping-to-user-part
        ping-from-user-part
        li-trust-me                     disabled
        in-manipulationid
        out-manipulationid
        manipulation-string
        manipulation-pattern
        p-asserted-id
        trunk-group
        max-register-sustain-rate       0
        early-media-allow
        invalidate-registrations        disabled
        rfc2833-mode                    none
        rfc2833-payload                 0
        codec-policy
        enforcement-profile
        refer-call-transfer             disabled
        refer-notify-provisional        none
        reuse-connections               NONE
        tcp-keepalive                   none
        tcp-reconn-interval             0
        max-register-burst-rate         0
        register-burst-window           0
        sip-profile
        sip-isup-profile
        kpml-interworking               inherit
        last-modified-by                admin@172.41.1.2
        last-modified-date              2013-07-03 11:47:20
sip-config
        state                           enabled
        operation-mode                  dialog
        dialog-transparency             enabled
        home-realm-id                   Core
        egress-realm-id
        nat-mode                        None
        registrar-domain                *
        registrar-host                  *
        registrar-port                  5060
        register-service-route          always
        init-timer                      500
        max-timer                       4000
        trans-expire                    32
        invite-expire                   180
        inactive-dynamic-conn           32
        enforcement-profile
        pac-method
        pac-interval                    10
        pac-strategy                    PropDist
        pac-load-weight                 1
        pac-session-weight              1
        pac-route-weight                1
        pac-callid-lifetime             600
        pac-user-lifetime               3600
        red-sip-port                    1988
        red-max-trans                   10000
        red-sync-start-time             5000
        red-sync-comp-time              1000
        add-reason-header               disabled
        sip-message-len                 4096
        enum-sag-match                  disabled
        extra-method-stats              disabled
        rph-feature                     disabled
        nsep-user-sessions-rate         0
        nsep-sa-sessions-rate           0
```



```
        registration-cache-limit       0
        register-use-to-for-lp         disabled
        refer-src-routing              disabled
        add-ucid-header                disabled
        proxy-sub-events
        allow-pani-for-trusted-only    disabled
        pass-gruu-contact              disabled
        sag-lookup-on-redirect         disabled
        set-disconnect-time-on-bye     disabled
        last-modified-by               admin@console
        last-modified-date             2013-05-13 15:05:35
sip-interface
        state                          enabled
        realm-id                       Access
        description
        sip-port
                address                        192.168.11.101
                port                           5060
                transport-protocol             UDP
                tls-profile
                multi-home-addrs
                allow-anonymous                registered
                ims-aka-profile
        carriers
        trans-expire                   0
        invite-expire                  0
        max-redirect-contacts          0
        proxy-mode
        redirect-action
        contact-mode                   none
        nat-traversal                  none
        nat-interval                   30
        tcp-nat-interval               90
        registration-caching           enabled
        min-reg-expire                 300
        registration-interval          3600
        route-to-registrar             enabled
        secured-network                disabled
        teluri-scheme                  disabled
        uri-fqdn-domain
        trust-mode                     all
        max-nat-interval               3600
        nat-int-increment              10
        nat-test-increment             30
        sip-dynamic-hnt                disabled
        stop-recurse                   401,407
        port-map-start                 0
        port-map-end                   0
        in-manipulationid
        out-manipulationid
        manipulation-string
        manipulation-pattern
        sip-ims-feature                disabled
        subscribe-reg-event            disabled
        operator-identifier
        anonymous-priority             none
        max-incoming-conns             0
        per-src-ip-max-incoming-conns  0
        inactive-conn-timeout          0
        untrusted-conn-timeout         0
        network-id
        ext-policy-server
        default-location-string
        charging-vector-mode           pass
```



```
        charging-function-address-mode pass
        ccf-address
        ecf-address
        term-tgrp-mode                 none
        implicit-service-route         disabled
        rfc2833-payload                101
        rfc2833-mode                   transparent
        constraint-name
        response-map
        local-response-map
        ims-aka-feature                disabled
        enforcement-profile
        route-unauthorized-calls
        tcp-keepalive                  none
        add-sdp-invite                 disabled
        add-sdp-profiles
        sip-profile
        sip-isup-profile
        tcp-conn-dereg                 0
        register-keep-alive            none
        kpml-interworking              disabled
        tunnel-name                    sipTunnel
        last-modified-by               admin@10.1.31.40
        last-modified-date             2013-05-10 12:06:29
sip-interface
        state                          enabled
        realm-id                       Core
        description
        sip-port
                address                        192.168.12.100
                port                           5060
                transport-protocol             UDP
                tls-profile
                multi-home-addrs
                allow-anonymous                agents-only
                ims-aka-profile
        carriers
        trans-expire                   0
        invite-expire                  0
        max-redirect-contacts          0
        proxy-mode
        redirect-action
        contact-mode                   none
        nat-traversal                  none
        nat-interval                   30
        tcp-nat-interval               90
        registration-caching           disabled
        min-reg-expire                 300
        registration-interval          3600
        route-to-registrar             disabled
        secured-network                disabled
        teluri-scheme                  disabled
        uri-fqdn-domain
        trust-mode                     all
        max-nat-interval               3600
        nat-int-increment              10
        nat-test-increment             30
        sip-dynamic-hnt                disabled
        stop-recurse                   401,407
        port-map-start                 0
        port-map-end                   0
        in-manipulationid
        out-manipulationid
        manipulation-string
```



```
        manipulation-pattern
        sip-ims-feature                 disabled
        subscribe-reg-event             disabled
        operator-identifier
        anonymous-priority              none
        max-incoming-conns              0
        per-src-ip-max-incoming-conns   0
        inactive-conn-timeout           0
        untrusted-conn-timeout          0
        network-id
        ext-policy-server
        default-location-string
        charging-vector-mode            pass
        charging-function-address-mode pass
        ccf-address
        ecf-address
        term-tgrp-mode                  none
        implicit-service-route          disabled
        rfc2833-payload                 101
        rfc2833-mode                    transparent
        constraint-name
        response-map
        local-response-map
        ims-aka-feature                 disabled
        enforcement-profile
        route-unauthorized-calls
        tcp-keepalive                   none
        add-sdp-invite                  disabled
        add-sdp-profiles
        sip-profile
        sip-isup-profile
        tcp-conn-dereg                  0
        register-keep-alive             none
        kpml-interworking               disabled
        tunnel-name
        last-modified-by                admin@console
        last-modified-date              2013-05-23 16:05:29
steering-pool
        ip-address                      172.16.0.201
        start-port                      49152
        end-port                        65535
        realm-id                        Access
        network-interface
        last-modified-by                admin@172.41.1.2
        last-modified-date              2013-07-03 14:28:18
steering-pool
        ip-address                      192.168.12.100
        start-port                      49152
        end-port                        65535
        realm-id                        Core
        network-interface
        last-modified-by                admin@console
        last-modified-date              2013-05-29 13:36:56
system-config
        hostname
        description                     BCP Access SBC
        location
        mib-system-contact
        mib-system-name
        mib-system-location
        snmp-enabled                    enabled
        enable-snmp-auth-traps          disabled
        enable-snmp-syslog-notify       disabled
        enable-snmp-monitor-traps       disabled
```



```
        enable-env-monitor-traps       disabled
        snmp-syslog-his-table-length   1
        snmp-syslog-level              WARNING
        system-log-level               WARNING
        process-log-level              NOTICE
        process-log-ip-address         0.0.0.0
        process-log-port               0
        collect
                sample-interval                5
                push-interval                  15
                boot-state                     disabled
                start-time                     now
                end-time                       never
                red-collect-state              disabled
                red-max-trans                  1000
                red-sync-start-time            5000
                red-sync-comp-time             1000
                push-success-trap-state        disabled
        call-trace                     disabled
        internal-trace                 disabled
        log-filter                     all
        default-gateway                172.40.0.1
        restart                        enabled
        exceptions
        telnet-timeout                 0
        console-timeout                0
        remote-control                 enabled
        cli-audit-trail                enabled
        link-redundancy-state          disabled
        source-routing                 disabled
        cli-more                       disabled
        terminal-height                24
        debug-timeout                  0
        trap-event-lifetime            0
        default-v6-gateway             ::
        ipv6-signaling-mtu             1500
        ipv4-signaling-mtu             1500
        cleanup-time-of-day            00:00
        snmp-engine-id-suffix
        snmp-agent-mode                v1v2
        comm-monitor
                state                          disabled
                qos-enable                     enabled
                sbc-grp-id                     0
                tls-profile
task done
```

