Oracle® Communications Subscriber-Aware Load Balancer

Essentials Guide

Release L-CX1.5.0

Formerly Net-Net Session-aware Load Balancer

November 2013

ORACLE®

Contents

About This Guide

Overview

Version L-CX1.5.0 provides an updated release of the Oracle Communications Subscriber-Aware Load Balancer (SLB). This guide describes that release.

Audience

This guide is written for network administrators and architects, and provides information about the SLB configuration. For information on configuration and operation of SLB cluster members refer to the 6.3.x, 6.4.x, or 7.1.2 documentation sets. Each of these Session Border Control (SBC) releases support SLB cluster membership, in addition to the full complement of other SBC functionality.

Related Documentation

Document Name	Document Description
Net-Net 4500 Hardware Installation Guide (Software Versions 6.3.x, 6.4.x, or 7.1.2)	Contains information about the components and installation of Acme Packet hardware.
Net-Net Configuration Guide (Software Versions 6.3.x, 6.4.x, or 7.1.2)	Contains information about the administration and configuration of the Acme Packet software.
Net-Net ACLI Reference Guide (Software Versions 6.3.x, 6.4.x, or 7.1.2)	Contains explanations of how to use the ACLI – provides alphabetical listings and descriptions of all ACLI commands and configuration parameters.
Net-Net Maintenance and Troubleshooting Guide (Software Versions 6.3.x, 6.4.x, or 7.1.2)	Contains information about Security Gateway logs, performance announcements, system management, inventory management, upgrades, working with configurations, and managing backups and archives.
Net-Net MIB Reference Guide (Software Versions 6.3.x, 6.4.x, or 7.1.2)	Contains information about Management Information Base (MIBs), Acme Packet's enterprise MIBs, general trap information, including specific details about standard traps and enterprise traps, Simple Network Management Protocol (SNMP) GET query information (including standard and enterprise SNMP GET query names, object identifier names and numbers, and descriptions), examples of scalar and table objects.

Document Revision History

This section contains a revision history for this document.

Date	Revision Number	Description
November 19, 2013		Initial release of software version LC-X 1.5.0

Subscriber-Aware Load Balancer Configuration

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 SBC's capacity: they will either use a DNS-based distribution mechanism, or they will pre-provision end-points 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 end-points.

The Subscriber-Aware Load Balancer (SLB) addresses the need for scaling a network edge to millions of end-points. Designed as a standalone system (an Acme Packet 4500 SBC) capable of supporting up to two million users (where a "user" is defined as a unique source IP address), and flexibly deployable into existing network topologies, the SLB aggregates signaling from large user populations to reduce the edge reachability problem by an order of magnitude.

> **Note:** The Enhanced Traffic Controller (ETC) Network Interface Unit (NIU) is not supported for SLB operations.

Functional Overview

The 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 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 new capabilities. First, the SBC supports the 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 an SLB-initiated *rebalance* operation, as described in "Rebalancing" on page 2. A group of SBCs, with the above-listed capabilities, that receive signaling traffic from the SLB is referred to as a *cluster*.

The IP-in-IP encapsulation technique provides SLB transparency to the terminating SBC. That is, when an SBC receives an encapsulated packet via the SLB, it can discard the outer encapsulation leaving behind an *identical* packet as transmitted originally by the end-point. Visibility into the actual packet transmitted by the end-point is necessary to provide certain services in the SBC (for example, hosted NAT traversal, session-agent matching, and so on). A secondary goal achieved by using this encapsulation technique is that it provides a disassociation function between an SBC's connected network and its SIP reachability. That is, an SBC can be assigned any IP address it wants from a network topology standpoint, yet still process SIP packets as though it were logically situated elsewhere at Layer 5. In a larger sense, the *physicality* of the SBC is no longer important; like-configured, logically identical SBCs can be spread all over the globe.

Balancing and Rebalancing

The SLB performs two primary functions as the front-end to an SBC cluster: balancing traffic and rebalancing traffic. There are several key distinctions, which are described in the following two sections.

Balancing

Balancing is defined as the distribution of new end-points (IP addresses) among the members of the SBC cluster. The SLB balances traffic based upon its configured policies (refer to "Load Balancer Policy Configuration" on page 13 for policy description and details), or, in the absence of configured policies, with a default round-robin procedure. Load balancer policies provide a flexible means of directing traffic to appropriate groups of SBCs. As initial packets arrive at the SLB from unknown (previously unseen) end-points, they are passed to a resident software process that consults its policy engine to determine an appropriate destination (clustered SBC) for each end-point. Regardless of the distribution algorithm, policy-based or round-robin, the SLB chooses an SBC from among all equally weighted candidates, giving preference to those with the lowest current *occupancy rate*, defined as the number of end-points already present on that system relative to its maximum end-point capacity.

Even though each clustered SBC regularly reports CPU date to the SLB, the SBC's CPU utilization is not factored into the preference of one SBC over another. Rather, a SBC whose CPU utilization rate exceeds its load limit threshold (by default, 90%) is excluded from the list of candidates. For example, assuming that both SBCs are licensed for the same number of sessions, an SBC with a CPU load of 89% and a current occupancy of 10,000 end-points will have equal footing with an SBC with a CPU load of 10% and a current occupancy of 10,000 end-points. But an SBC with a CPU load of 90% and an occupancy of 0 end-points will never receive new assignments from the SLB, until its CPU utilization rate falls below the 90% threshold.

Rebalancing

Rebalancing, as opposed to balancing, is taking some number of existing users from functioning SBCs and redistributing these existing users between current cluster members. Rebalancing can be automatically scheduled when a new SBC joins an existing cluster, or immediately invoked with the Acme Packet Command Line Interface (ACLI). When an SBC exits a cluster, whatever the reason, all of its end-points are invalidated on the SLB and those users are essentially *balanced* when they revisit the SLB.

A new SBC joins an existing cluster by initiating the establishment of an IP-in-IP tunnel between itself and the SLB. During an initial handshake the SBC designates which SLB service port or ports it is prepared to support. If there are existing SBCs supporting these designated service ports, the SLB instructs some or all of these SBCs to divest themselves of a specified number of end-points. The SLB calculates the number of divested end-points based upon the overall occupancy of that service relative to the SBC's contribution to that occupancy. Existing cluster members not advertising support for service ports designated by the new cluster member are excluded from the rebalance queue.

The SLB sequences through eligible cluster member one at a time, using Cluster Control Protocol (CCP) messaging to request nomination and removal of eligible end-points. The SBC replies with a CCP response that lists candidate end-points. The SLB removes existing forwarding rules associated with those end-points, and repeats the CCP request/response process until the cluster member divests itself of the specified number of end-points.

When the divested end-points re-engage with the SLB (upon their next scheduled registration refresh, for example), the SLB lacks a forwarding rule that maps them to a specific SBC. Consequently, the message is passed up to the software processes running on the SLB's host, which chooses a new destination for that user – presumably, the new cluster member that has the most available capacity.

The cluster member, after being requested to nominate end-points for rebalancing, uses several criteria for choosing the most attractive candidates. As part of its standard SIP processing performed by SBCs, the cluster member is aware of the expiry times for all of the entries in its SIP registration cache. Therefore, the cluster member can predict with a high degree of accuracy when any given end-point will be signaling back into the cluster. As the forwarding rules on the cluster member are triggered by end-point messages, the cluster member considers a user whose registration entry is due to expire shortly an attractive candidate for rebalance. Note, however, that in many cases it is not prudent to nominate users whose SIP registration cache entries are due to expire *immediately*, as this can cause a race condition between the CCP response and the SIP REGISTER message from the end-point to the SIP registration function. To avoid this potential dilemma, cluster members are equipped with the ability to skip ahead to candidates whose expiry is not immediate.

Further, each cluster member categorizes the end-points stored in its cache based upon a priority value that is determined via the SLB's distribution policy (see "Distribution Policy Configuration" on page 16 for more details). It nominates end-points from its lowest priority buckets first.

Finally, the SLB does not rebalance an active SIP end-point — an end-point engaged in a phone conversation.

After removing end-points from the first cluster member, the SLB moves to the next cluster member in the rebalance queue. and uses the same CPP request/response exchange to remove additional end-points. The same procedure repeats for additional cluster members until the SLB attains the target number of divested end-points.

When the divested end-points re-engage with the SLB (upon their next scheduled registration refresh, for example), the SLB lacks a forwarding rule that maps them to a specific cluster member. Consequently, the message is passed up to the software processes running on the SLB's host, which chooses a new destination for that user – presumably, the new cluster member that has the most available capacity.

IPv4/IPv6 Dual Stack

While major carriers are proceeding toward a pure IPv6 network for next generation services, current practicalities require the continued support of IPv4 handsets and other devices. As a result, the SLB provides support for single non-channelized physical interfaces that support both IPv4 and IPv6 ingress and egress on the same network interface.

Support for the dual stack interface requires no new additional configuration elements, and is provided by the proper configuration of the following elements:

Physical Interfaces

Refer to the section entitled *Physical Interfaces: Net-Net 4500 SBC* in the *System Configuration* chapter of the latest version of the Net-Net 4000 ACLI Configuration Guide (release versions 6.3.x, 6.4.x, or 7.1.2).

IPv4 Network Interfaces

Refer to the section entitled *About Your Net-Net 3800/4500 and IPv6* in the *System Configuration* chapter of the latest version of the Net-Net 4000 ACLI Configuration Guide (release versions 6.3.x, 6.4.x, or 7.1.2).

IPv6 Network Interfaces

Refer to the sections entitled *Configuring Network Interfaces: Net-Net 4500 SBC*, *Licensing*, *Globally Enabling IPv6*, *IPv6 Address Configuration*, *IPv6 Default Gateway*, and *Network Interfaces and IPv6* in the *System Configuration* chapter of the latest version of the Net-Net 4000 ACLI Configuration Guide (release versions 6.3.x, 6.4.x, or 7.1.2).

IPv4 SIP Interfaces

Refer to the section entitled *SIP Interfaces* in the *SIP Signalling Services* chapter of the latest version of the Net-Net 4000 ACLI Configuration Guide (release versions 6.3.x, 6.4.x, or 7.1.2) (release versions 6.3.x, 6.4.x, or 7.1.2).

IPv6 SIP Interfaces

Refer to the section entitled *SIP Interfaces* in the *SIP Signalling Services* chapter of the latest version of the Net-Net 4000 ACLI Configuration Guide.

IPv4 Realms

Refer to the section entitled *Configuring Realms* in the *Realms* chapter of the latest version of the Net-Net 4000 ACLI Configuration Guide (release versions 6.3.x, 6.4.x, or 7.1.2).

IPv6 SIP Interfaces

Refer to the section entitled *Configuring Realms* in the *Realms* chapter of the latest version of the Net-Net 4000 ACLI Configuration Guide (release versions 6.3.x, 6.4.x, or 7.1.2).

SLB Configuration

This section explains how to configure functionality specific to the SLB; it does not include configuration steps for functions that it shares in common with its corresponding SBCs (for example, *system-config*, *phy-interface*, *network-interface*, and so on). For information about general SBC configuration, refer to the appropriate documentation as listed in *About This Guide*.

SLB configuration is quite simple; aside from basic network connectivity, the service interfaces, and the distribution policy, much of the configuration is learned dynamically from the SBCs that comprise the cluster.

SLB Tunnel Configuration

The SLB sends and receives signaling messages to and from clustered SBCs through an IP-in-IP tunnel. The SLB requires one tunnel per interface.

Use the following procedure to perform required SLB-side tunnel configuration. Completion of tunnel configuration is accomplished on the clustered SBCs as described in "SBC Tunnel Configuration" on page 22.

1. From superuser mode, use the following ACLI command sequence to access *tunnel-config* configuration mode. While in this mode, you partially configure the *tunnel-config* configuration element.

```
alesmith# configure terminal
alesmith(configure)# system
alesmith(system)# network-interface
alesmith(network-interface)# tunnel-config
alesmith(tunnel-config)# ?

local-ip-address             tunnel local IP address
port                         tunnel local & remote control ports
protocol                     tunnel control transport protocol
tls-profile                  tunnel control TLS profile
select                       select tunnel to edit
no                           delete tunnel
show                         show tunnel
done                         write tunnel information
exit                         return to previous menu

alesmith(tunnel-config)#
```

2. Use the **local-ip-address** parameter to specify the IP address at the SLB end of the tunnel.

   As the terminus for all tunnels from the clustered SBCs — and never the tunnel originator — only the local address is configured on the SLB.

   Note: This address also supports the exchange of CCP messages.

```
alesmith(tunnel-config)# local-address 182.16.204.210
alesmith(tunnel-config)#
```

3. Use the **port** parameter to specify the port used to send and receive CCP messages.

```
alesmith(tunnel-config)# port 4444
alesmith(tunnel-config)#
```

4. Use the **protocol** parameter to specify the transport protocol used in support of cluster control messages.

   Supported protocols are UDP (the recommended default), TCP, or TLS.

   Note: CCP messages are exchanged quite frequently, and the overhead associated with encrypting and decrypting these messages is significant. Selection of TLS as the Transport Layer protocol degrades system performance.

   ```
   alesmith(tunnel-config)# protocol UDP
   alesmith(tunnel-config)#
   ```

5. If TLS is the selected transport protocol, use the **tls-profile** parameter to select the existing TLS profile that identifies the cryptographic resources used to secure the TLS connection.

   ```
   alesmith(tunnel-config)# tls-profile TLS-LB
   alesmith(tunnel-config)#
   ```

6. Use **done**, **exit**, and **verify-config** to complete configuration of this *tunnel-config* configuration element.

7. Repeat Steps 1 through 6 to configure additional *tunnel-config* configuration elements.

Sample SLB Tunnel Configuration

The following formatted extract from **show running-config** ACLI output shows a sample tunnel configuration.

```
tunnel-config
local-ip-address     182.16.204.210
port                 4444
protocol             UDP
tls-profile
last-modified-by     admin@console
last-modified-date   2013-11-07 18:49:04
```

Cluster Configuration

The *cluster-config* configuration element manages basic SLB interaction with clustered SBCs — it contains a set of global parameters that define the management of the RFC 2003 IP-in-IP tunnels that connect the SLB to clustered SBCs, and the details of rebalance operations. In addition, *cluster-config* provides for the creation of a list of *service interfaces* (signaling addresses) that are advertised to end-points comprising the user access population.

Use the following procedure to perform required *cluster-config* configuration.

1. From superuser mode, use the following ACLI command sequence to access *cluster-config* configuration mode. While in this mode, you configure the *cluster-config* configuration element.

   ```
   alesmith# configure terminal
   alesmith(configure)# session-router
   alesmith(session-router)# cluster-config
   alesmith(cluster-config)# ?
   ```

```
state                          cluster control state
log-level                      configure log level
auto-rebalance                 Auto-rebalance cluster on new SD
                               availability
source-rebalance-threshold     Percentage of advertised registration
                               capacity
dest-rebalance-threshold       Percentage of advertised registration
                               capacity
dest-rebalance-max             Percentage of advertised registration
                               capacity
tunnel-check-interval          How often an SD's tunnels are checked
tunnel-fail-interval           Time for which no messages have been
                               received
rebalance-request-delay        Delay between subsequent rebalance
                               requests
session-multiplier             ratio of users (end-points to sessions)
atom-limit-divisor             ratio of atoms (e.g. contacts to
                               end-points)
rebalance-skip-ahead           Skip end-points refreshing sooner than
rebalance-max-refresh          Skip end-points refreshing later than
ignore-tgt-svcs-on-rebalance   When selecting source SDs during
                               rebalancing
rebalance-del-app-entries      Delete Application end-point Data
inactive-sd-limit              Duration no SD control messages received
                               (seconds)
red-port                       redundant mgcp sync port
red-max-trans                  max redundant transactions to keep
red-sync-start-time            redundant sync start timeout
red-sync-comp-time             redundant sync complete timeout
service-ports                  configure service ports
select                         select cluster config
no                             delete cluster config
show                           show cluster config
done                           save cluster config information
exit                           return to previous menu

alesmith(cluster-config)#
```

2. Use the **state** parameter to enable or disable the SLB software.

   The default setting, *enabled*, enables SLB functionality; *disabled* renders the SLB inoperable.

```
alesmith(cluster-config)# state enabled
alesmith(cluster-config)#
```

3. Use the **log-level** parameter to specify the contents of the SLB log.

   Log messages are listed below in descending order of severity.

emergency	— the most severe
critical	
major (error)	
minor (error)	
warning	
notice	
info	— (default) the least severe
trace	— (test/debug, not used in production environments)
debug	— (test/debug, not used in production environments)
detail	— (test/debug, not used in production environments)

In the absence of an explicitly configured value, **log-level** defaults to *critical*, meaning that log messages with a severity of *critical* or greater (*emergency*) are written to the SLB log.

```
alesmith(cluster-config)# log-level critical
alesmith(cluster-config)#
```

4. Use the **auto-rebalance** parameter to specify SLB behavior when a new SBC joins an existing cluster.

   With this parameter *enabled*, the default setting, the SLB redistributes end-points among cluster members when a new member joins the cluster. Refer to“Rebalancing” on page 2 for operational details.

   With this parameter *disabled*, the alternate setting, pre-existing SBCs retain their end-point populations, and the SLB directs all new end-points to the newly active SBC until that SBC reaches maximum occupancy.

   ```
   alesmith(cluster-config)# auto-rebalance enabled
   alesmith(cluster-config)#
   ```

5. If **auto-rebalance** is set to *enabled*, use the **source-rebalance-threshold** and **dest-rebalance-threshold** parameters to specify threshold settings that identify existing cluster SBCs as either end-point sources or end-point destinations during the rebalance operation. Use the **dest-rebalance-max** parameter to specify the occupancy for the new cluster member. Refer to “Balancing” on page 2 for details on occupancy and its calculation.

   If **auto-rebalance** is set to *disabled*, these three parameters can be ignored.

   Parameter values are numeric percentages within the range 0 through 100.

   source-rebalance-threshold specifies the minimum occupancy percent that identifies a clustered SBC as a source of end-points during a rebalance operation. For example, using the default value of *50* (percent), any clustered SBC with an occupancy rate of 50% or more sheds end-points during a rebalance. The SLB assigns these end-points to the new cluster member.

   dest-rebalance-threshold specifies the maximum occupancy percent that identifies a clustered SBC as a destination for end-points during a rebalance operation. Note that the default setting of *0* (percent), ensures that no pre-existing SBC gains end-points during a rebalance.

   dest-rebalance-max specifies the maximum occupancy percent that the SLB transfers to the new cluster member during a rebalance operation. The default setting is *80* (percent). Should this threshold value be attained, the SLB distributes remaining end-points to those SBCs identified as end-point destinations by their **dest-rebalance-threshold** settings.

   ```
   alesmith(cluster-config)# source-rebalance-threshold 50
   alesmith(cluster-config)# dest-rebalance-threshold 40
   alesmith(cluster-config)# dest-rebalance-max 75
   ```

6. If **auto-rebalance** is set to *enabled*, you can optionally use four additional parameters to fine-tune rebalance operational details.

   If **auto-rebalance** is set to *disabled*, these four parameters can be ignored.

   rebalance-request-delay specifies the interval (in milliseconds) between end-point request messages sent from the SLB to a clustered SBC. As explained in “Rebalancing” on page 2 these messages request a list of end-points that will be redistributed from the SBC to a new cluster member.

   By default, this parameter is set to 500 milliseconds.

Setting this parameter to a higher value results in longer times for the completion of rebalancing; however longer durations provide more time for cluster member processsing of SIP traffic.

rebalance-skip-ahead restricts the target set of SBC end-points registration eligible for rebalancing to those whose re-registration is not imminent — that is, the registration is not scheduled within the number of milliseconds specified by the parameter setting. Setting this parameter to a non-zero value mitigates against the possibility of a race condition precipitated by a simultaneous end-point removal generated by the SBC and the arrival of end-point signalling on an SLB service port. The default setting (*0* milliseconds) effectively makes the entire SBC end-point set eligible for rebalancing.

rebalance-max-refresh restricts the target set of SBC end-points eligible for rebalancing to those whose re-registration is no further in the future than the time period (milliseconds) specified by this parameter— for example, assuming a parameter value of 6000, the target end-point set is restricted to those whose re-registration is scheduled within the next 6 seconds.

Because a re-balancing operation necessarily introduces a small window of unreachability for re-balanced end-points, this parameter provides users with some degree of control over the period of time that a re-balanced end-point may be unreachable.

The default setting (*0* milliseconds) effectively makes the entire SBC end-point set eligible for rebalancing.

rebalance-del-app-entries specifies when cached SIP entries for rebalanced end-points are removed from the clustered SBC. The default setting (*disabled*) specifies that cached entries are retained after a rebalance operation, and subsequently removed from the cache by standard time-out procedures. When set to *enabled*, this parameter specifies that the SBC removes cached registration entries at the completion of the rebalance operation.

```
alesmith(cluster-config)# rebalance-request-delay 750
alesmith(cluster-config)# rebalance-skip-ahead 100
alesmith(cluster-config)# rebalance-max-refresh 1000
alesmith(cluster-config)# rebalance-del-app-entries enabled
```

7. Three parameters, **tunnel-fail-interval**, **tunnel-check-interval,** and **inactive-sd-limit** maintain and monitor the IP-in-IP tunnels established between the SLB and clustered SBCs.

   tunnel-fail-interval specifies the interval (in milliseconds) between periodic *keepalive* messages sent from a clustered SBC to the SLB. If the SLB fails to receive a *keepalive* message within the specified period, it flags the tunnel as *dead*. By default, this parameter is set to *10000* milliseconds.

   tunnel-check-interval specifies the interval (in milliseconds) between SLB tunnel audits. During a tunnel audit, the SLB checks the status of each tunnel and removes all tunnels flagged as *dead*. If all of a cluster member's tunnels are removed. the SLB places that cluster member in an *out-of-service* state. By default, this parameter is set to *15000* milliseconds.

   If you change default settings for either parameter, ensure that the setting for **tunnel-check-interval** is greater than the **tunnel-fail-interval** setting.

   inactive-sd-limit specifies the maximum silent interval (defined as the absence of heartbeat traffic from any tunnel) seconds) before the SLB flags a cluster member as *dead*, and removes that SBC from the cluster. By default, this parameter is set to *1800* seconds (30 minutes). supported values are integers within the range 0 through 31556926 (365 days).

```
alesmith(cluster-config)# tunnel-fail-interval 10000
alesmith(cluster-config)# tunnel-check-interval 15000
alesmith(cluster-config)# inactive-sd-limit 900
```

8. Use the **session-multiplier** and **atom-limit-divisor** parameters to specify optional, user-configurable numeric factors used in *occupancy* and *occupancy rate* calculations.

   session-multiplier provides a factor that when multiplied by an SBC's licensed session limit, determines the maximum number of end-points that the SBC can support (that is, its maximum *occupancy*).

   The default setting is *10*; valid settings include any integer values within the range 1 through 100.

   Using the default setting, an SBC licensed for 32,000 concurrent sessions has a maximum theoretical occupancy of 320,000 end-points.

   atom-limit-divisor provides another factor that can be used in *occupancy* and *occupancy percent* calculations. By default, occupancy calculations are based on end-points (IP addresses), and do not take into account the fact that the same IP address can represent multiple users.

   The default setting is *1*, which assumes a conservative 1-to-1 correlation between end-points and users; valid settings include any integer values within the range 1 through 1000.

   Note: The SLB initially calculates a tentative maximum occupancy value, expressed as a number of end-point addresses, for each clustered SBC. SLB calculations are based upon the licensed capacity of each cluster member, and the values assigned to the **session-multiplier** and **atom-limit-divisor** parameters. After calculating the tentative maximum occupancy value, the SLB compares this value to the value of the **registration-cache-limit** parameter as defined on the clustered SBC. If the value of **registration-cache-limit** is either 0, or greater than the tentative maximum occupancy value, the calculated value is retained as the occupancy ceiling. However, if the **registration-cache-limit** value is greater than 0, but less than the tentative calculation, the value of **registration-cache-limit** is used as the occupancy ceiling.

   Once an SBC has reached its maximum number of end-points, the SLB removes it from the load balancing algorithm. These parameter settings should be changed only after careful examination of network conditions and behavior.

```
alesmith(cluster-config)# session-multiplier 10
alesmith(cluster-config)# atom-limit-divisor 1
```

9. The **ignore-tgt-svc-on-rebalance** parameter is not currently supported, and can be safely ignored.

10. Retain default settings for the **red-port**, **red-max-trans**, **red-sync-start-time**, and **red-sync-comp-time** parameters.

11. Use **done**, **exit**, and **verify-config** to complete cluster configuration.

Sample Cluster Configuration

The following formatted extract from **show running-config** ACLI output shows a sample cluster configuration.

```
cluster-config
state                          enabled
log-level                      CRITICAL
auto-rebalance                 enabled
source-rebalance-threshold     50
dest-rebalance-threshold       40
dest-rebalance-max             75
tunnel-check-interval          750
tunnel-fail-interval           10000
rebalance-request-delay        500
session-multiplier             4
rebalance-skip-ahead           0
rebalance-max-refresh          0
ignore-tgt-svcs-on-rebalance   disabled
atom-limit-divisor             1000
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
last-modified-by               admin@console
last-modified-date             2013-11-07 18:49:04
```

Service Ports Configuration

A service port is essentially a SIP port monitored by the SLB for incoming signaling from the user population. For virtually all network topologies, multiple service ports are expected on a typical SLB configuration. A *service-port* is a multiple instance configuration element; for each service port advertised to the access network(s), at least one *service-port* configuration element must be configured.

Use the following procedure to perform required *service-ports* configuration.

1. From superuser mode, use the following ACLI command sequence to access *service-port* configuration mode. While in this mode, you configure one or more *service-port* configuration elements.

   ```
   alesmith# configure terminal
   alesmith(configure)# session-router
   alesmith(session-router)# cluster-config
   alesmith(cluster-config)# service-ports
   alesmith(service-port)# ?
   ```

```
address                        IP address
port                           port (default: 5060)
protocol                       transport protocol
network-interface              network interface for service port
select                         select cluster config
no                             delete cluster config
show                           show cluster config
done                           save cluster config information
exit                           return to previous menu

alesmith(service-port)#
```

2. Use the required **address** parameter to specify the IPv4 or IPv6 address of this service port.

   ```
   alesmith(service-port)# address 10.0.0.1
   alesmith(service-port)#
   ```

3. Use the **port** parameter to specify the port monitored by the SLB for incoming signaling messages.

   In the absence of an explicitly configured port, the SLB provides a default value of 5060 (the registered SIP port).

   Allowable values are integers within the range 0 through 65535.

   ```
   alesmith(service-port)# port 5060
   alesmith(service-port)#
   ```

4. Use the **protocol** parameter to choose the transport protocol.

   Supported settings are *UDP* (the recommended default), All, and TCP.

   If the value, *All*, is chosen, the value of the **port** parameter if forced to *0*, a special value designating all ports from 1 through 65535.

   ```
   alesmith(service-port)# protocol udp
   alesmith(service-port)#
   ```

5. Use the required **network-interface** parameter to identify the SLB network interface that supports this service port.

   ```
   alesmith(service-port)# network-interface M00:0
   alesmith(service-port)#
   ```

6. Use **done**, **exit**, and **verify-config** to complete configuration of this *service-port* configuration element.

7. Repeat Steps 1 through 6 to configure additional *service-port* configuration elements.

Sample Service Port Configuration

The following formatted extract from **show running-config** ACLI output shows sample service port configurations.

```
service-port
address                         192.169.203.83
port                            5060
protocol                        UDP
network-interface               M00:0
last-modified-by                admin@console
last-modified-date              2013-11-07 18:49:04

service-port
address                         192.169.203.83
port                            5050
protocol                        TCP
network-interface               M00:0
last-modified-by                admin@console
last-modified-date              2013-11-07 18:49:04
```

Load Balancer Policy Configuration

The *lbp-config* configuration element manages the SLB end-point table. It also creates and manages a list of *service interfaces* (signaling addresses) that are advertised to end-points comprising the user access population.

Use the following procedure to perform required *lbp-config* configuration.

1. From superuser mode, use the following ACLI command sequence to access *lbp-config* configuration mode. While in this mode, you configure the *lbp-config* configuration element.

```
alesmith# configure terminal
alesmith(configure)# session-router
alesmith(session-router)# lbp-config
alesmith(lbp-config)#?

state                                lbp state
log-level                            configure log level
red-port                             lbp redundant sync port: 0 to
                                     disable and 2000 to enable
red-max-trans                        maximum redundancy transactions
                                     to keep on active
red-sync-start-time                  timeout for transitioning
                                     from standby to active
red-sync-comp-time                   sync request timeout after initial
                                     sync completion
untrusted-grace-period               Untrusted grace period
max-untrusted-percentage             Maximum untrusted end-points
                                     percentage
max-untrusted-upper-threshold        Maximum untrusted end-points
                                     upper threshold
max-untrusted-lower-threshold        Maximum untrusted end-points
                                     upper threshold
end-point-capacity-upper-threshold   end-point capacity upper
                                     threshold
end-point-capacity-lower-threshold   end-point capacity lower
                                     threshold
```

```
options                          optional features/parameters
select                           select lbp config
no                               delete lbp config
show                             show lbp config
done                             save lbp config information
exit                             return to previous menu

alesmith(lbp-config)#
```

2. Use the **state** parameter to enable or disable the SLB software.

   The default setting, *enabled*, enables SLB functionality; *disabled* renders the SLB inoperable.

   ```
   alesmith(lbp-config)# state enable
   alesmith(lbp-config)#
   ```

3. Use the **log-level** parameter to specify the contents of the SLB log.

   Log messages are listed below in descending order of severity.

emergency	— the most severe
critical	
major (error)	
minor (error)	
warning	
notice	
info	— (default) the least severe
trace	— (test/debug, not used in production environments)
debug	— (test/debug, not used in production environments)
detail	— (test/debug, not used in production environments)

   In the absence of an explicitly configured value, **log-level** defaults to *critical*, meaning that log messages with a severity of *critical* or greater (*emergency*) are written to the LBP log.

   ```
   alesmith(lbp-config)# log-level critical
   alesmith(lbp-config)#
   ```

4. Use the **untrusted-grace-period**, **max-untrusted-percentage**, **max-untrusted-upper-threshold**, and **max-untrusted-lower-threshold** parameters to implement percentage-based management and monitoring of untrusted end-points in the SLB end-point database. Management and monitoring of untrusted end-points is instrumental in detecting and responding to Denial-of-Service (DOS) attacks aimed at the SLB.

   untrusted-grace-period specifies the maximum time, in seconds, that a forwarding rule is retained by the SLB before it is confirmed with a promotion message from the SBC which received the untrusted end-point. Refer to “Balancing” on page 2 for message details.

   In the absence of an explicitly assigned value, the SLB provides a default setting of *30* (seconds).

   If this time period elapses without a promotion message arriving to confirm this user, the SLB deletes the entry.

   Setting this parameter to 0 allows untrusted/unconfirmed entries to exist indefinitely without aging out.

   max-untrusted-percentage specifies the percentage of the overall end-point population that is reserved for untrusted users.

   The default setting is *20* (percent); supported values are integers within the range 1 through 100.

This percentage is applied to the overall remaining occupancy of the SLB after trusted (confirmed) users are accounted for. For example, when empty, the SLB holds two million forwarding rules; assuming the default setting, at most 400,000 rules are reserved for untrusted rules. By the time one million users have been promoted, 20% of the remaining space means that up to 200,000 entries can be used for untrusted users.

max-untrusted-upper-threshold specifies a threshold level at which the SLB (1) raises an alarm, and (2) issues an SNMP trap reporting an excessive number of untrusted end-points within the entire end-point population.

This parameter, which has a default setting of *80* (percent), is calculated as a percent of **max-untrusted-percentage**. For example, assuming default settings for both parameters, the SLB raises an alarm and issues an SNMP trap when the percentage of untrusted end-points attains 16%.

max-untrusted-lower-threshold specifies a threshold level at which the SLB (1) clears the existing untrusted end-point alarm, and (2) issues an SNMP trap reporting alarm clearance.

This parameter, which has a default setting of *70* (percent), is calculated as a percent of **max-untrusted-percentage**. For example, assuming default settings for both parameters, the SLB clears an alarm and issues an SNMP trap when the percentage of untrusted end-points falls to 14%.

```
alesmith(lbp-config)# untrusted-grace-period 30
alesmith(lbp-config)# max-untrusted-percentage 20
alesmith(lbp-config)# max-untrusted-upper-threshold 80
alesmith(lbp-config)# max-untrusted-lower-threshold 70
alesmith(lbp-config)#
```

5. Use the **end-point-capacity-upper-threshold** and **end-point-capacity-lower-threshold** parameters to implement license-based management and monitoring of the SLB end-point counts.

   end-point-capacity-upper-threshold specifies a threshold level at which the SLB (1) raises an alarm, and (2) issues an SNMP trap reporting an excessive number of active end-points.

   This parameter, which has a default setting of *80* (percent), is calculated as a percentage of the end-points allowed by the installed SLB license.

   end-point-capacity-lower-threshold specifies a threshold level at which the SLB (1) clears the existing end-point alarm, and (2) issues an SNMP trap reporting alarm clearance.

   This parameter, which has a default setting of *70* (percent), is calculated as a percentage of the end-points allowed by the installed SLB license.

```
alesmith(lbp-config)# end-point-capacity-upper-threshold 80
alesmith(lbp-config)# end-point-capacity-lower-threshold 70
alesmith(lbp-config)#
```

6. Use **done**, **exit**, and **verify-config** to complete configuration of this *load-balancer-policy* configuration element.

Sample Load Balancer Policy Configuration

The following formatted extract from **show running-config** ACLI output shows a sample load balancer policy configuration.

```
lbp-config
state                                enabled
log-level                            NOTICE
untrusted-grace-period               30
max-untrusted-percentage             20
max-untrusted-upper-threshold        80
max-untrusted-lower-threshold        70
end-point-capacity-upper-threshold    80
end-point-capacity-lower-threshold    70
red-port                             0
red-max-trans                        500000
red-sync-start-time                  5000
red-sync-comp-time                   1000
last-modified-by                     admin@console
last-modified-date                   2013-11-07 18:49:04
```

Distribution Policy Configuration

Distributing end-points equitably among the cluster members is the primary function of the SLB. The *lp-config* configuration element allows you to control the method of the SLB's distribution based on matching criteria. Using inbound packet matching criteria, you can control the assignment of users to SBCs. Matching is done by data available up to and including the transport layer of the packet: source IP address and port, destination IP address and port, and transport protocol. The IP addresses and ports may or may not include bit masks as well.

Conceptually, the load balancer policy table, with sample data, looks akin to the following.

Source IP/Mask	Source Port/ Mask	Destination IP/Mask	Destination Port/Mask	Transport Protocol Require-ments (list)	Realm Identi-fiers (list)
192.168.7.22/32	0/0	10.0.0.1/32	5060/16		West
192.168.1.0/24	0/0	10.0.0.1/32	5060/16	UDP, TCP	North, South, West
192.168.0.0/16	0/0	10.0.0.1/32	5060/16	UDP, TCP	East, West
0.0.0.0/0	0/0	0.0.0.0/0	0/0		

Policies are matched using a longest prefix match algorithm; the most specific policy is selected when comparing policies to received packets. One and only one policy is chosen per packet; if the next hops in that route are all unavailable, the next best route is not consulted (instead, the *default policy* may be consulted – see below). This is different than the local-policy behavior on the SBC.

Within each policy you may configure multiple next hops, where each next hop is a named group of SBCs. In the sample policy table, this is indicated in the second policy with a source IP range of 192.168.1.0/24. The realm identifier list for this policy indicates North, South, West. Each of these realm identifiers represents a collection of zero or more SBCs, in SBC parlance these are roughly analogous to *session-agent groups*. Each of these realm identifiers is also assigned a *priority* (a value between 1 and 31, with 31 representing the highest priority) in the configuration, and the SLB sorts the possible

destinations with the highest priority first. Upon receipt of a packet matching a policy with multiple configured realm identifiers, the SLB gives preference to SBCs from the realm identifier with the highest priority. Should no SBCs be available in that priority level (due to saturation, unavailability, and so on.) the SLB moves on to investigate the next priority level, and so on. Should no SBCs become available after traversing the entire list of all SBCs within each priority level, the SBC either drops the packet or attempt to use the *default policy*.

The bottom row of the sample table shows this implicit, last resort default policy. When enabled, the SLB reverts to the default policy when all of the potential next hop realms referenced in the end-point's distribution rule are unavailable. In that event, the default policy attempts to locate a clustered SBC that advertises support for the service-interface that the packet arrived on. The realm is not considered when matching to the default policy. If such an SBC is found, the SLB forwards the packet to that SBC; if such an SBC is not found, the SLB drops the packet.

It is not necessary to configure the default policy — it is simply intended as a catchall policy, and may be used when all that is required is a simple round-robin balancing scheme based on simple metrics (for example, CPU utilization and number of registrations currently hosted by an SBC). If no policies are configured on the SLB, the default policy is used. The default realm is implied in the above table as "*" and is enabled by default for policy records.

Use the following procedure to perform required *lb-config* configuration.

1. From superuser mode, use the following ACLI command sequence to access *lb-config* configuration mode. While in this mode, you configure the distribution rules used to implement policy-based load balancing on the SLB.

```
alesmith# configure terminal
alesmith(configure)# session-router
alesmith(session-router)# lb-policy
alesmith(lb-policy)#?

state                              lb policy state
default-realm                      use default realm
description                        load balancer policy description
protocols                          list of protocols
lb-realms                          list of realms
                                        name
                                        priority
source-addr                        source ip address
destination-addr                   destination ip address
select                             select lb policy
no                                 delete lb policy
show                               show lb policy
done                               save lb policy information
exit                               return to previous menu

alesmith(lb-policy)#
```

2. Use the **state** parameter to enable or disable this distribution rule.

   The default setting, *enabled*, enables the distribution rule; *disabled* disables the rule.

```
alesmith(lb-policy)# state enabled
alesmith(lb-policy)#
```

3. Use the **default-realm** parameter to enable or disable the default distribution policy.

   The default setting, *enabled*, enables the default policy; *disabled* disables the policy.

   With **default-realm** enabled, the SLB provides a best-effort delivery model if the next-hop realms listed in this distribution rule are unavailable. With **default-realm** disabled, the orphaned packet is dropped.

   ```
   alesmith(lb-policy)# default-realm enabled
   alesmith(lb-policy)#
   ```

4. Optionally use the **description** parameter to provide a description of this distribution rule.

   ```
   alesmith(lb-policy)# description “Local traffic to Los Angeles site”
   alesmith(lb-policy)#
   ```

5. Use the **protocols** parameter to construct a list of protocols that must be supported by this distribution rule.

   ```
   alesmith(lb-policy)# protocols “udp tcp”
   alesmith(lb-policy)#
   ```

6. Use either the **source-addr** parameter or the **destination-address** parameter to specify matching criteria for this distribution rule.

   Use the **source-addr** parameter to specify source-address-based matching criteria.

   Packets whose source IP addresses match the criteria specified by this parameter are subject to this distribution rule.

   ```
   alesmith(lb-policy)# source-addr 10.0.0.1
   alesmith(lb-policy)#
   ```

   matches any port on the specified IP source address

   ```
   alesmith(lb-policy)# source-addr 10.0.0.1:5060
   alesmith(lb-policy)#
   ```

   matches the specified IP source address:port pair

   ```
   alesmith(lb-policy)# source-addr 10.0.0.1/24
   alesmith(lb-policy)#
   ```

   matches any IP source address, any port on the 10.0.0.x subnet

   ```
   alesmith(lb-policy)# source-addr 10.0.0.240/28:5060
   alesmith(lb-policy)#
   ```

   matches IP source addresses 10.0.0.240:5060 through 10.0.0.255:5060

   Use the **destination-addr** parameter to specify destination-address-based matching criteria.

   Packets whose destination IP addresses match the criteria specified by this parameter are subject to this distribution rule.

   ```
   alesmith(lb-policy)# destination-addr 10.0.0.1
   alesmith(lb-policy)#
   ```

   matches any port on the specified IP destination address

   ```
   alesmith(lb-policy)# destination-addr 10.0.0.1:5060
   alesmith(lb-policy)#
   ```

   matches the specified IP destination address:port pair

```
alesmith(lb-policy)# destination-addr 10.0.0.1/24
alesmith(lb-policy)#
```

matches any IP destination address, any port on the 10.0.0.x subnet

```
alesmith(lb-policy)# destination-addr 10.0.0.240/28:5060
alesmith(lb-policy)#
```

matches destination IP addresses 10.0.0.240:5060 through 10.0.0.255:5060

7. Use the **lb-realms** parameter to access *lb-realm* configuration mode.

   While in *lb-realm* configuration mode you identify one or more SBCs eligible to receive traffic that matches this distribution rule.

```
alesmith(lb-policy)# lb-realms
alesmith(lb-realm)#

name                                realm name (string identifier)
priority                            priority (range 1-31)
select                              select a lb realm to edit
no                                  delete selected lb realm
show                                show lb realm information
done                                write lb realm information
exit                                return to previous menu

alesmith(lb-realm)#
```

8. Use the **name** parameter to identify the realm.

   As previously discussed, the *name* field is roughly analogous to an SBC session-agent group. SBCs configured to communicate within a cluster hosted by an SLB advertise offered services to the SLB. These services (for example, SIP support) exist in realms, whose names are sent to the SLB as part of the SBC advertisement. The SLB, upon receipt of these advertisements, joins each SBC into one or more *realm identifier groups* based upon the realm name(s) the SBC has offered up. The **name** command of the *lb-realm* configuration element matches this distribution rule to a supporting SBC that has offered that realm name for cluster membership.

```
alesmith(lb-realm)# name LosAngeles
alesmith(lb-realm)#
```

9. Use the **priority** parameter to specify the realm priority.

   Priority is expressed as an integer value within the range 0 to 31 — the higher the integer, the greater the priority.

   The default value, 0, specifies use of the default routing policy, and should not be used when policy-based distribution is enabled.

   Priority values are considered when multiple SBCs offer the same service to matched packets.

```
alesmith(lb-realm)# priority 31
alesmith(lb-realm)#
```

10. Use **done**, **exit**, and **verify-config** to complete configuration of this *lb-realm* configuration element.

11. To specify other eligible SBCs, repeat Steps 7 through 10. For example,

```
alesmith(lb-config)# lb-realms
alesmith(lb-realm)# name LasVegas
alesmith(lb-realm)# priority 25
alesmith(lb-realm)# done
alesmith(lb-realm)# exit
alesmith(lb-realm)# verify-config
```

12. Use **done**, **exit**, and **verify-config** to complete configuration of this distribution rule.
13. To specify additional distribution rules, repeat Steps 1 through 12 as often as necessary.

Sample Distribution Rule Configurations

The following formatted extract from **show running-config** ACLI output shows sample distribution rule configurations.

```
lb-policy
state               enabled
default-realm       enabled
description
protocols           TCP

       lb-realm
       name         Realm192p1
       priority     10

source-addr         1.1.0.0/16
destination-addr    0.0.0.0/0
last-modified-by    admin@console
last-modified-date  2013-11-07 18:58:10

lb-policy
state               enabled
default-realm       enabled
description
protocols           TCP

       lb-realm
       name         Realm192p1
       priority     7

source-addr         1.20.0.0/16
destination-addr    0.0.0.0/0
last-modified-by    admin@console
last-modified-date  2013-11-07 19:01:01

lb-policy
state               enabled
default-realm       enabled
description
protocols           TCP
```

```
        lb-realm
        name         Realm192p1
        priority     5

source-addr         1.120.0.0/16
destination-addr    0.0.0.0/0
last-modified-by    admin@console
last-modified-date  2013-11-07 19:00:49

lb-policy
state               enabled
default-realm       enabled
description
protocols           TCP

        lb-realm
        name         Realm192p1
        priority     3
```

Forced Rebalance

The **notify ccd rebalance** ACLI command initiates an immediate forced rebalance operation. A forced rebalance operation is identical to the one described in “Rebalancing” on page 2.

notify ccd rebalance [cancel [sd-name]]

```
alesmith# notify ccd rebalance
alesmith#
```

initiates the forced rebalance by calculating drop counts for each eligible cluster member, and then requesting drops from the first cluster member in the rebalance queue.

```
alesmith# notify ccd rebalance cancel
alesmith#
```

terminates the forced rebalance.

```
alesmith# notify ccd rebalance cancel~sam
alesmith#
```

terminates the forced rebalance for a specified cluster member. Note the use of tilde special character, which forces the SLB to do a substring match of the following string against all cluster member names. Assuming a cluster member samadams@172.30.68.31 — that cluster member removes itself from the rebalance queue, if it has not yet removed end-points, or ceases end-point removal and exits the queue if it is currently doing so.

notify ccd drop <sd-name> (<realm> <number> | <number>)

```
alesmith# notify drop ~sam boston 100
alesmith#
```

instructs the target cluster member to drop 100 end-points from the boston realm

```
alesmith# notify drop ~sam * 100
alesmith#
```

using the * special character instructs the target cluster member to drop 100 end-points from all realms

```
alesmith# notify drop ~sam 100
alesmith#
```

instructs the target cluster member to drop 100 end-points without regard for realm

SBC Configuration

This section describes the configuration necessary to allow an SBC to join a cluster. Configuration is simplified to allow for an easy and seamless migration from a deployed standalone SBC to a deployed clustered SBC. There are only two places where new configuration is required: in the *network-interface* configuration element, where tunnel information is defined; and in the signaling application's interface, (the *sip-interface* configuration element).

SBC Tunnel Configuration

Configuring the properties of the IP-in-IP tunnel on the SBC is a matter of configuring the local address and specifying transport layer and application layer protocol support.

The following example uses a tunnel named *sipSignaling*, which was initially and partially configured on the SLB. Note in the following configuration that the **remote-address** parameter is not specified — that value having been previously set with the **local-address** parameter on the SLB. The complementary configuration performed on the SBC enables tunnel establishment between the SBC and the SLB.

1. From superuser mode, use the following ACLI command sequence to access *tunnel-config* configuration mode. While in this mode, you perform required SBC tunnel configuration.

```
westy# configure terminal
westy(configure)# system
westy(system)# network-interface
westy(network-interface)# tunnel-config
westy(tunnel-config)# ?

name                  tunnel name
local-ip-address      tunnel local IP address
remote-mac-address    tunnel remote mac address
remote-ip-address     tunnel remote IP address
application           application protocol for this tunnel
port                  tunnel local & remote control ports
protocol              tunnel control transport protocol
tls-profile           tunnel control TLS profile
select                select tunnel to edit
no                    delete tunnel
show                  show tunnel
done                  write tunnel information
exit                  return to previous menu

westy(tunnel-config)#
```

2. Use the **name** command to provide a unique identifier for this tunnel instance.

```
westy(tunnel-config)# name sipSignaling
westy(tunnel-config)#
```

3. Use the **local-address** parameter to specify the IP address at the SBC end of the tunnel.

   Note: This address also supports the exchange of CCP messages.

   ```
   westy(tunnel-config)# local-address 1.1.1.100
   westy(tunnel-config)#
   ```

4. Ignore the **remote-mac-address** parameter which is not required for tunnel configuration.

5. Use the **port** parameter to specify the port used to send and receive cluster control messages.

   ```
   westy(tunnel-config)# port 4444
   westy(tunnel-config)#
   ```

6. Use the **protocol** parameter to specify the transport protocol used in support of cluster control messages.

   Supported transport protocols are UDP (the recommended default), TCP, or TLS.

   Note: Cluster control messages are exchanged quite frequently, and the overhead associated with encrypting and decrypting these messages is significant. Selection of TLS as the Transport Layer protocol degrades system performance.

   ```
   westy(tunnel-config)# protocol UDP
   westy(tunnel-config)#
   ```

7. If TLS is the selected transport protocol, use the **tls-profile** parameter to select the existing TLS profile that identifies the cryptographic resources used to secure the TLS connection.

   TLS usage is not recommended!

   ```
   westy(tunnel-config)# tls-profile TLS-ClusterMember
   westy(tunnel-config)#
   ```

8. Use the **application** parameter to specify the application protocol supported by this tunnel.

   Specify the SIP protocol.

   ```
   westy(tunnel-config)# application SIP
   westy(tunnel-config)#
   ```

9. Use **done**, **exit**, and **verify-config** to complete configuration of this *tunnel-config* configuration element.

10. Repeat Steps 1 through 9 to complete tunnel configuration on other SIP interfaces as required.

Sample SBC Tunnel Configuration

The following formatted extract from **show running-config** ACLI output shows a sample SBC (cluster member) configuration.

```
tunnel-config
name                          one
local-ip-address              1.1.1.100
remote-mac-address
remote-ip-address
port                          4444
protocol                      UDP
tls-profile
application                   SIP
last-modified-by              admin@console
last-modified-date            2013-11-10 23:24:15
```

SIP Configuration

In a traditional SBC configuration the IP address assigned to a *sip-port* configuration element is contained within the address space defined by the network interface netmask. This is not be the case for clustered SBCs. Rather, the IP address assigned to the sip-port is identical to the address of an SLB service-port advertised on the access network. The process of encapsulating the packets between the SLB and SBC masks the fact that the IP address the SBC expects to receive IP packets on is different than the Layer 5 address the SBC expects the SIP address on.

Consistency of realm identification is vital to successful and predictable policy-based load balancing. Take particular care to ensure that the **realm-id** of the *sip-interface* configuration element mirrors the **lb-realm** assignments made while configuring distribution rules. See "Distribution Policy Configuration" on page 16.

In the following configuration example, the **realm-id** is *LosAngeles*. This SBC, when booted, will detect that it is a member of an SLB cluster and register the service port 10.0.0.1:5060/UDP as the realm *LosAngeles* with the SLB. The SLB will automatically create the SBC group *LosAngeles* (if it doesn't exist) or join the SBC to the group *LosAngeles* (if it is not the first to advertise *LosAngeles*). Policy statements that direct packets to *LosAngeles* now consider this SBC as a potential destination, assuming the address:port/protocol also are consistent with the policy's matching criteria.

This technique allows you to configure the same IP:port/protocol on multiple SBCs, with different realm-id labels, to indicate priority of one SBC or group of SBCs over another. As an example, consider several SBCs geographically situated together with the label *LosAngeles*, and several other SBCs geographically situated elsewhere with the label *NewYork*, all with the identical SIP interface and SIP port configuration. A policy can be easily defined to give preference to a source subnet of users in California to the *LosAngeles* member SBCs, with *NewYork* as a second priority. This provides flexibility in network design without undue burden in the configuration: SBCs' tagged with the same realm name are joined in dynamically created SBC groups by the SLB, with no explicit configuration required on the SLB whatsoever.

1. From superuser mode, use the following ACLI command sequence to access *sip-interface* configuration mode. While in this mode, you verify the **realm-id** and assign the newly created IP-in-IP tunnel to a SIP interface.

```
westy# configure terminal
westy(configure)# session-router
westy(session-router)# sip-interface
westy(sip-interface)# select
<realm-id>: LosAngeles
1: LosAngeles  172.192.1.15:5060

selection: 1
westy(sip-interface)# show
sip-interface
        state                   enabled
        realm-id                LosAngeles
        ...
        ...
        ...
westy(sip-interface)#
```

2. Use the **tunnel-name** parameter to assign the IP-in-IP tunnel to the current SIP interface.

```
westy(sip-interface)# tunnel-name sipSignaling
westy(sip-interface)# ?
```

3. Use the **sip-port** command to move to *sip-port* configuration mode.

```
westy(sip-interface)# sip-port
westy(sip-port)# ?

address            IP Address
port               port (default: 5060)
transport-protocol transport protocol
tls-profile        the profile name
allow-anonymous    allowed requests from SIP realm
ims-aka-profile    ims-aka profile name
select             select a sip port to edit
no                 delete a selected sip port
show               show sip port information
done               write sip port information
exit               return to previous menu
westy(sip-port)#
```

4. Use the **address**, **port**, and **transport-protocol** parameters to mirror the address of an existing SLB service port.

```
westy(sip-port)# address 10.0.0.1
westy(sip-port)# port 5060
westy(sip-port)# transport-protocol udp
westy(sip-port)#
```

5. Use **done**, **exit**, and **verify-config** to complete configuration of this *sip-port* configuration element.

6. Repeat Steps 1 through 5 as necessary to verify **realm-id**s, assign IP-in-IP tunnels, and create mirrored service ports on additional SIP interfaces.

Online/Offline Configuration

The **set-system-state** ACLI command provides the ability to temporarily place a clustered SBC in the *offline* state. The offline setting puts the SBC into a state where it is powered on and available only for administrative purposes.

The transition to the offline state is graceful in that existing calls are not affected by the state transition. The SBC informs the SLB of the impending status change via a CCP message. Upon receiving such a message, the SLB ceases to forward new end-points to the SBC, and places the SBC in the *Shutdown* state. The SBC, for its part, enters a state that results in the rejection of any incoming out-of-dialog SIP requests. Eventually all calls compete, registrations expire and are removed by the SLB, and returning end-points are allotted to active SBCs.

Use the **set-system-state offline** ACLI command to place an SBC in the offline state.

```
westy# set-system-state offline
Are you sure you want to bring the system offline? [y/n]?: y
Setting system state to going-offline, process will complete when all
current calls have completed
westy#
```

Note: An SBC in the offline state plays no role in a balance or rebalance operation.

In a similar fashion use the **set-system-state online** ACLI command to place an SBC in the online state.

```
westy# set-system-state online
Are you sure you want to bring the system online? [y/n]?: y
Setting system state to online
westy#
```

SLB/Cluster Management & Diagnostics

SLB Statistics

The SLB provides the operator with a full set of statistical data for troubleshooting and diagnostic purposes. This section describes current statistical outputs and defines displayed values. It is important to become familiar with the data and the collection process when opening trouble tickets as service personnel will rely upon this information to assist you in diagnosing hardware, software, and/or network issues.

show balancer

The **show balancer** command is the root of all statistical data pertinent to SLB operation. Below is a list of valid arguments, which are described in further detail in the following sections:

```
alesmith# show balancer ?

end-points  show session load balancer end-points
members     show session load balancer cluster member summary
metrics     show load balancer metrics
realms      show load balancer realms
tunnels     show session load balancer statistics
statistics  show session load balancer IP-in-IP tunnel info
alesmith#
```

show balancer end-points

The **show balancer end-points** command displays a full list of all IP-to-SBC mappings resident in the SLB. As the SLB can hold up to two million entries, the output of this command can and will grow very large, and extreme caution should be exercised when executing this command on a heavily trafficked SLB system.

```
alesmith# show balancer end-points
IP address       Port Index    Address  Flags      SBC Handle
--------------- ----- -------- -------- ---------- ----------
14.0.134.236     5060 001e81e7 001e81e7 c0000000           94
14.0.134.232     5060 001e81eb 001e81eb c0000000           93
14.0.134.228     5060 001e81ec 001e81ec c0000000           92
14.0.134.224     5060 001e81ea 001e81ea c0000000           98
14.0.134.220     5060 001e81ee 001e81ee c0000000           97
14.0.134.216     5060 001e81ef 001e81ef c0000000           96
14.0.134.212     5060 001e81f1 001e81f1 c0000000           95
14.0.134.208     5060 001e81f4 001e81f4 c0000000           94
14.0.134.204     5060 001e81f0 001e81f0 c0000000           93
14.0.134.200     5060 001e81f3 001e81f3 c0000000           92
14.0.134.196     5060 001e81e7 001e81e7 c0000000           98
14.0.134.192     5060 001e81f9 001e81f9 c0000000           97
14.0.134.188     5060 001e81f6 001e81f6 c0000000           96
alesmith#
```

The table provided by **show balancer end-points** displays every end-point mapping. In the above example, note that IP addresses in the 14.0.134.0/24 space are being distributed among a number of SBCs. The *IP address* and *Port* columns pinpoint a specific end-point. The *Index*, *Address*, and *Flags* columns contain SLB internal reference identifiers for locating that specific end-point in memory. The *SBC Handle* column identifies which SBC serves that end-point; use the **show balancer members** command to display a mapping of SBC names to SBC handles.

The **notify lbp save-stats** command saves the entire end-point table to flash in a compressed (*xz*) format.

> **Note: notify lbp save-stats** may take several minutes to run, and should not be executed while the SLB is performing initial ramp-up of end-point traffic, nor while the SLB is actively rebalancing.

You can use optional command arguments to filter/restrict command output.

show balancer end-points <*ip-address*> restricts the display to one end-point.

For example:

show balancer end-points address 14.0.134.232

displays data for the specified IP end-point

show balancer end-points <*ip-address*>/<*:port_num*>> restricts the display to a specific port on a specific IP address.

For example:

show balancer end-points address 14.0.134.232:5060

displays data for port 5060 on the specified end-point

show balancer end-points <*ip-address*>/<*bit-mask-len*> restricts the display to a contiguous range of end-point addresses.

For example:

show balancer end-points address 14.0.134.0/24

displays data for the 14.0.134.0 subnet

show balancer end-points address 14.0.134.240/28

displays data for end-point addresses 14.0.134.240 through 14.0.134.255

show balancer end-points <*ip-address*>/<*bit-mask-len*><*:port_num*>>

displays data for a specific port on a contiguous range of end-point addresses

show balancer members

The **show balancer members** command provides a list of all SBCs that have registered with the SLB.

```
alesmith# show balancer members
SBC Name  Source Address    Destination Address  S/Port/Vlan end-points
--- ----  ---------------   -------------------  ----------- -----------
87 csbc3  00:08:25:a1:ee:f9 00:08:25:a1:fe:39 0  1/0/0       76592
88 csbc1a 00:08:25:a1:ee:f9 00:08:25:a1:ef:3f 0  1/0/0       72399
...
...
...

              max end-points: 2000000
    max untrusted end-points: 292777
          current end-points: 526142
current untrusted end-points: 186074
               current SBCs: 7
alesmith#
```

SBC contains the SBC handle, an internal shorthand that identifies a specific SBC. The **show balancer members** command provides a handle-to-hostname mapping.

Name contains the SBC hostname. Standalone SBCs are displayed as *hostname@IP address*, and highly available SBCs (*csbc1a* in the above display) are displayed as *hostname*.

Source IP contains the local (SLB) IPv4 or, as displayed above, IPv6 tunnel address.

Destination IP contains the remote (SBC) IPv4 or, as displayed above, IPv6 tunnel address.

Slot, *Port*, and *Vlan* identify the local interface that supports the SLB-to-SBC tunnel.

end-points contains the number of end-point-SBC associations that the SLB created for each specific SBC,

max end-points contains the licensed capacity of the SLB.

max untrusted end-points contains the maximum allowed number of untrusted end-points.

current end-points contains the current number of end-points, trusted and untrusted

current untrusted end-points contains the current number of untrusted end-points.

show balancer metrics

The **show balancer metrics** command displays a comparison between the number of *local end-points* (that is, the associations between source addresses and each SBC) and the number of *remote end-points* (that is, what the SBC reports to the SLB as the number of end-points it has received via the tunneled interface). Note that in the example output below those two numbers are the same; this is true if and only if there are no users in the access network that have multiple phone lines sourced from the same IP address. Were that the case, the number of remote end-points would be higher than the number of local end-points.

This table is populated with the data received in the periodic heartbeats from the SBC to the SLB. As these heartbeats are somewhat infrequent (every two seconds by default), the data in this table should only be considered accurate within two seconds.

```
alesmith# show balancer metrics
                              local  remote                  max
SBC Name                      epts   epts    max reg   CPU   CPU
--- ------------------------- ------ ------- -------- ----  ----
 93 magichat@172.30.68.34     0      0       480000    2.7  90.0
 94 westy                     0      0       480000    2.7  90.0
 95 samadams@172.30.68.33     0      0       480000    2.8  90.0
 96 bass@172.30.68.35         0      0       480000    4.3  90.0
 97 sixtus@172.30.68.36       0      0       480000    2.9  90.0
 98 newcastle@172.30.68.37    0      0       480000    2.9  90.0
 99 guiness@172.30.68.33      0      0       480000    3.6  90.0
alesmith#
```

SBC contains the SBC handle.

Name contains the SBC hostname.

max reg contains the maximum number of end-points the SLB will send to this specific SBC. Its value is derived from the product of the **session-multiplier** parameter in the *cluster-config* configuration element and the SBC's licensed session capacity. The SBC passes this value to the SLB during the SBC's registration process into the cluster.

CPU contains the last received information on the CPU percentage from this SBC.

max CPU contains the threshold percentage at which the SBC is removed from consideration for the assignment of new end-points. The default value is 90%, and may be changed on an SBC by setting the load-limit value as a SIP configuration option.

show balancer realms

The **show balancer realms** command displays a composite list of realms that all member SBCs have registered with the SLB.

```
alesmith# show balancer realms
Realm        SBC Tunnel Name                    ref count end-points
----------- --- ------ ------------------------ --------- ---------
access       99   4092 newcastle@172.30.68.37           1     53535
access       98   4091 magichat@172.30.68.34            1     53535
access       97   4090 augustiner@172.30.68.41          1     53535
access       94   4086 bass@172.30.68.35                1     53535
access       93   4085 westy@172.30.68.42               1     53535
access       92   4084 sixtus@172.30.68.36              1     53535
access       96   4089 guiness@172.30.68.33             1     53535
net-13       99   4088 samadams@172.30.68.31            1     62550
net-13       91   4087 stbernie@172.30.68.43            1     62550
alesmith#
```

In this example, seven of the nine SBCs have registered the realm *access* and two have registered the realm *net-13*. The total number of end-points for each of these services is indicated in the rightmost column. *ref count* is reserved for future use.

show balancer statistics

The **show balancer statistics** command displays statistical output pertinent to low-level events on the SLB. The contents and output of this command are subject to change, and will be documented in a subsequent document release.

```
alesmith# shower-balancer-statistics
LBP not initialized drops          0
max capacity reached drops         2
end-point SBC mismatch errors      0
end-point table read errors        0
Tx packet failed count             0
service not found count            0
duplicate ept packet drops         0
msgq drops                         0
forwarded duplicates               0
policy miss                    40508
realm miss                         0
throttle drops                     0
throttle skips                     0
throttle policy skips              0
total packets processed        40508
end-points removed                 0
EPT delete errors                  0
end-points added                   0
EPT add errors                     0
EPT update errors                  0
group not found                    0
LBP agent not found                0
invalid end-point                  0
insert error                       0
untrusted dropped                  0
untrusted age outs               171
packets dropped in balance         0
packets dropped by standby         0
-----------------------------
trusted end-points (EPT db)        0
untrusted end-points (EPT db)    337
-----------------------------
total trusted end-points           0
total untrusted end-points       337
total end-points                 337
alesmith#
```

show balancer tunnels

When implemented on the SLB, the **show balancer tunnels** command generates a list of data for each tunnel between the SLB and its clustered SBCs. It includes the tunnel source and destination addresses, as well as an internal switch ID (swid) for this tunnel.

```
alesmith# show balancer tunnels
4090(32752)::
outer src ip = 182.16.203.83
outer dst ip = 182.16.203.87
slot/port/vlan = 1/0/0
trusted swid/pipe = 40/1213
untrusted swid/pipe = 41/1212
host swid/pipe = 42/1211
  service: 192.169.203.83:5050 [Realm192p1] protocols: 6/21588
  service: 192.169.203.83:5060 [Realm192p1] protocols: 17/21588

4091(32753)::
outer src ip = 182.16.203.83
outer dst ip = 182.16.203.86
slot/port/vlan = 0/1/0
trusted swid/pipe = 36/1216
untrusted swid/pipe = 37/1215
host swid/pipe = 38/1214
  service: 192.169.203.83:5060 [Realm192p1] protocols: 17/21588

alesmith#
```

Use the **error** argument for error reporting and troubleshooting.

```
alesmith# show balancer tunnels errors
src ip 182.16.203.83 / dst ip 182.16.203.87 / slot 0 / port 1 / vlan 0:
  IP:Port:  192.169.203.83:5050
  Proto Encaps Errors Decaps Errors
  ----- ------------- -------------
  6                 0             0
  IP:Port:  192.169.203.83:5060
  Proto Encaps Errors Decaps Errors
  ----- ------------- -------------
  17                0             0

src ip 182.16.203.83 / dst ip 182.16.203.86 / slot 0 / port 1 / vlan 0:
  IP:Port:  192.169.203.83:5060
  Proto Encaps Errors Decaps Errors
  ----- ------------- -------------
  17                0             0

unknown protocol:      0
do not fragment drops: 0
no matching tunnel:    0
service lookup failed: 0
IP frag msg failure:   0
mblk alloc failures:   0
unknown errors:        0
alesmith#
```

The **show balancer tunnels error** command can also be executed on an SBC cluster member. In this usage, the displayed data is restricted to errors between the specific cluster member and the SLB.

Use the **fragments** argument for information related to packet fragmentation/reassembly details.

```
alesmith# show balancer tunnels fragments
src ip 182.16.203.83 / dst ip 182.16.203.87 / slot 0 / port 1 / vlan 0:
  IP:Port:  192.169.203.83:5050
  Proto Encaps Fragmented Decaps Reassembled
  ----- ---------------- ------------------
  6                    0                 0
  IP:Port:  192.169.203.83:5060
  Proto Encaps Fragmented Decaps Reassembled
  ----- ---------------- ------------------
  17                   0                 0

src ip 182.16.203.83 / dst ip 182.16.203.86 / slot 0 / port 1 / vlan 0:
  IP:Port:  192.169.203.83:5060
  Proto Encaps Fragmented Decaps Reassembled
  ----- ---------------- ------------------
  17                   0                 0

MTU = 1500

packets reassembled:    0
fragments too small:    0
fragment len not mod 8: 0
too many fragments:     0
reassembly timeouts:    0
duplicate fragments:    0
mbuf alloc failed:      0
gap in reassembly:      0
IP hdr checksum errors: 0
overlapping frags:      0
packet too long:        0
multi-buffer frames:    0
unknown errors:         0
alesmith#
```

The **show balancer tunnels fragments** command can also be executed on an SBC cluster member. In this usage, the displayed data is restricted to fragmentation operations between the specific cluster member and the SLB.

Use the **statistics** argument for information related to packet counts.

```
alesmith# show balancer tunnels statistics
src ip 182.16.203.83 / dst ip 182.16.203.87 / slot 0 / port 1 / vlan 0:
  IP:Port:  192.169.203.83:5050
  Proto Encap Pkts Encap Octets Decap Pkts Decap Octets
  ----- ---------- ------------ ---------- ------------
  6              0            0          0            0
  IP:Port:  192.169.203.83:5060
  Proto Encap Pkts Encap Octets Decap Pkts Decap Octets
  ----- ---------- ------------ ---------- ------------
  17         48011     24213914          0            0

src ip 182.16.203.83 / dst ip 182.16.203.86 / slot 0 / port 1 / vlan 0:
  IP:Port:  192.169.203.83:5060
  Proto Encap Pkts Encap Octets Decap Pkts Decap Octets
  ----- ---------- ------------ ---------- ------------
  17         48017     24217918          0            0


alesmith#
```

The **show balancer tunnels statistics** command can also be executed on an SBC cluster member. In this usage, the displayed data is restricted to traffic counts between the specific cluster member and the SLB.

Cluster Control Protocol Statistics

The CCP provides the operator with a full set of statistical data for troubleshooting and diagnostic purposes.

show ccd

The **show ccd** command is the root of all statistical data pertinent to CCP operation. Below is a list of valid arguments, which are described in further detail in the following sections:

```
alesmith# show ccd ?

ccp         Cluster Control Protocol Stats
rebalance   Display rebalance queue
reset       Reset Stats
sds         Controlled SDs
stats       Cluster Control Stats
alesmith#
```

show ccd ccp

The **show ccd ccp** command displays aggregated data (that is, from all cluster members) about specific CCP operations.

```
alesmith# show ccd ccp
--------------------
M01:0
--------------------

Svc Add                Recent      Total  PerMax
===================    ======  =========  ======
Ops Recvd                   0          8       4
Op Replies Sent             0          8       4

                       -------- Received -------   ----------  Sent --------
Status Code            Recent      Total  PerMax   Recent      Total  PerMax
-----------            ------  ---------  ------   ------  ---------  ------
200 OK                      0          0       0        0          8       4

EP Del                 Recent      Total  PerMax
===================    ======  =========  ======
Ops Recvd                   0       3984    1013
Duplicate Ops               0         12       6
Op Replies Sent             0       3984    1013

                       -------- Received -------   ----------  Sent --------
Status Code            Recent      Total  PerMax   Recent      Total  PerMax
-----------            ------  ---------  ------   ------  ---------  ------
200 OK                      0          0       0        0       3984    1013

EP Promo               Recent      Total  PerMax
===================    ======  =========  ======
Ops Recvd                   0     115601    8187
Duplicate Ops               0        329      23
Op Replies Sent             0     115601    8187
```

```
                     -------- Received -------   ----------  Sent --------
Status Code          Recent      Total  PerMax   Recent      Total  PerMax
-----------          ------  ---------  ------   ------  ---------  -
200 OK                    0          0       0        0     115601    8187

Metrics              Recent      Total  PerMax
===================  ======  =========  ======
Ops Recvd                57      16330      75
Op Replies Sent          57      16331      75

                     -------- Received -------   ----------  Sent --------
Status Code          Recent      Total  PerMax   Recent      Total  PerMax
-----------          ------  ---------  ------   ------  ---------  -
200 OK                    0          0       0       57      16328      75
406 Not Accept            0          0       0        0          3       3

Prov Done            Recent      Total  PerMax
===================  ======  =========  ======
Ops Recvd                 0          6       3
Op Replies Sent           0          6       3

                     -------- Received -------   ----------  Sent --------
Status Code          Recent      Total  PerMax   Recent      Total  PerMax
-----------          ------  ---------  ------   ------  ---------  -
200 OK                    0          0       0        0          6       3

Going Down           Recent      Total  PerMax
===================  ======  =========  ======
Ops Recvd                 0          1       1
Op Replies Sent           0          1       1

                     -------- Received -------   ----------  Sent --------
Status Code          Recent      Total  PerMax   Recent      Total  PerMax
-----------          ------  ---------  ------   ------  ---------  -
200 OK                    0          0       0        0          1       1

Stop Down            Recent      Total  PerMax
===================  ======  =========  ======
Ops Recvd                 0          8       6
Op Replies Sent           0          9       6

                     -------- Received -------   ----------  Sent --------
Status Code          Recent      Total  PerMax   Recent      Total  PerMax
-----------          ------  ---------  ------   ------  ---------  -
200 OK                    0          0       0        0          6       3
406 Not Accept            0          0       0        0          3       3

alesmith#
```

Use the *hostname* argument to display data for a specific cluster member.

```
alesmith# show ccd ccp westy
--------------------
westy
--------------------

Svc Add                Recent      Total  PerMax
===================   ======  =========  ======
Ops Recvd                  0          1       1
Op Replies Sent            0          1       1


                      -------- Received ------- ---------- Sent --------
Status Code            Recent      Total  PerMax Recent      Total  PerMax
-----------            ------  ---------  ------   ------ ---------  ------
200 OK                     0          0       0 0          1       1

EP Del                 Recent      Total  PerMax
===================   ======  =========  ======
Ops Recvd                  0          1       1
Op Replies Sent            0          1       1

                      -------- Received ------- ---------- Sent --------
Status Code            Recent      Total  PerMax Recent      Total  PerMax
-----------            ------  ---------  ------   ------ ---------  ------
200 OK                     0          0       0 0          1       1

EP Promo               Recent      Total  PerMax
===================   ======  =========  ======
Ops Recvd                  0       1002     947
Op Replies Sent            0       1002     947

                      -------- Received ------- ---------- Sent --------
Status Code            Recent      Total  PerMax Recent      Total  PerMax
-----------            ------  ---------  ------   ------ ---------  ------
200 OK                     0          0       0 0       1002     947

Metrics                Recent      Total  PerMax
===================   ======  =========  ======
Ops Recvd                 20     167849      15
Op Replies Sent           20     167849      15

                      -------- Received ------- ---------- Sent --------
Status Code            Recent      Total  PerMax Recent      Total  PerMax
-----------            ------  ---------  ------   ------ ---------  ------
200 OK                     0          0       0 20     167849      15

Prov Done              Recent      Total  PerMax
===================   ======  =========  ======
Ops Recvd                  0          1       1
Op Replies Sent            0          1       1

                      -------- Received ------- ---------- Sent --------
Status Code            Recent      Total  PerMax Recent      Total  PerMax
-----------            ------  ---------  ------   ------ ---------  ------
200 OK                     0          0       0 0          1       1

Stop Down              Recent      Total  PerMax
===================   ======  =========  ======
Ops Recvd                  0          5       2
Op Replies Sent            0          5       2
```

```
                        -------- Received ------- ---------- Sent --------
Status Code             Recent      Total  PerMax Recent      Total  PerMax
-----------             ------  ---------  ------   ------ ---------  ------
200 OK                       0          0       0 0          5       2

alesmith#
```

show ccd sds

The **show ccd sds** command displays a table containing an overview of all of the data gleaned from the CCP from each SBC.

```
alesmith# show ccd sds
Session Director        Hdl  State     Tunnel Svcs Version     HW
----------------------- ---- --------- ------ ---- ---- ------ -----
augustiner@172.30.68.41   95 InService  2/2     2  6.2.0.30b8  SD3
bass@172.30.68.35         94 InService  2/2     2  6.2.0.30b8  SD3
guinness@172.30.68.33     96 InService  2/2     2  6.2.0.30b8  SD3
magichat@172.30.68.34     97 InService  2/2     2  6.2.0.30b8  SD3
newcastel@172.30.68.37    98 InService  2/2     2  6.2.0.30b8  SD3
samadams@172.30.68.31     99 InService  2/2     2  6.2.0.30b8  SD3
sixtus@172.30.68.36       92 InService  2/2     2  6.2.0.30b8  SD3
stbernie@172.30.68.43     91 InService  2/2     2  6.2.0.30b8  SD3
westy@172.30.68.42        93 InService  2/2     2  6.2.0.30b8  SD3
alesmith#
```

Session Director contains the hostname of the cluster SBCs that are connected to the SLB. As with all of the similar statistical output, standalone SBCs are displayed as *hostname@eth0 IP address*, and highly available SBCs are displayed as *hostname*.

Hdl contains the clustered SBC handle, an internal shorthand that identifies a specific cluster member. The **show balancer members** command provides a handle to hostname mapping.

State contains the current SBC state. Valid states are:

- *Init* — during initial handshaking with the SLB
- *InService* — healthy and operating normally
- *Rebalance* — during a cluster expansion/contraction operation
- *LostControl* — no longer communicating with the SLB

Tunnel contains the number of tunnels between the SBC and SLB.

Svcs contains the number of advertised services (protocols) that the SBC has negotiated with the SLB.

Version contains the software version running on that SBC.

HW identifies the hardware platform (in this case, SD3 identifies an Acme Packet Net-Net 4500 SBC).

LastPing is not currently used.

When issued with an optional *hostname* argument, the **show ccd sds** command provides a detailed report for the target hostname.

```
alesmith# show ccd sds bass
Session Director: bass@172.30.68.35
+-------------------------------------------------------------------
|State     : InService       Handle       : 0x5f
|Tunnels   : 1               ServicePorts : 1
|HW Type   : SD3             SW Version   : 6.2.0.30b8
|Last Ping : 1080ms          App Count    : 1
|
|Service:        App  SvcPorts Tunnels end-points DropCount
+--------------- ---- -------- ------- --------- ---------
access           SIP         1       1    285714         0
|
|#   Tunnel                              App  Handle     Svcs LastHB
+--- ---------------------------------- ---- ---------- ---- -----
|0   (1.1.1.100|1.1.1.15)               SIP  0xffb         1  375ms
|
|#   CPU   MAX   CurReg    RegLimit  CurSess   MaxSess
+--- ----- ----- --------- --------- --------- ---------
|0    4.1% 90.0%    285714         0      7336     64000
|     4.1% 90.0%    285714    960000      7736     64000
|
|Service Port                          App  Handle      TunNdx Avail
+------------------------------------- ---- ---------- ------ -----
|access::192.168.168.100:5060<17>      H248     513(1)      0   yes
alesmith#
```

SBC State

State — the current SBC state

Handle — the SBC handle

Tunnels — the current number of SBC tunnels

ServicePorts — the current number of SBC service ports

HW Type — the hardware platform (in this case, SD3 identifies an Acme Packet Net-Net 4500 SBC)

SW Version — the installed software revision level

Last Ping — the number of elapsed milliseconds, since a ping/keepalive was received from this SBC

App Count — the number of applications supported by the SBC

Services State

Service — the realm advertised by the SBC in the Service Port ID

App — the supported protocol: SIP

SVCPorts — the current number of service ports

Tunnels — the current number of tunnels

end-points — the cumulative number of end-points for this service

DropCount — the number of elements to drop when rebalancing this SBC

Tunnel State

#	—	the tunnel index (0 or 1)
Tunnel	—	the SLB and SBC tunnel IP address
App	—	the supported protocol: SIP
Handle	—	the handle for the tunnel
Svcs	—	the number if service ports supporting the tunnel
LastHB	—	the number of elapsed milliseconds since a heartbeat was received from the remote end of this tunnel

Tunnel Metrics

#	—	the tunnel number (0 or 1)
CPU	—	the current CPU utilization rate
Max	—	the maximum supported CPU utilization rate, if this value is exceeded, the tunnel implements a load limit algorithm
CurReg	—	the current number of registrations supported by the SBC
regLimit	—	the maximum number of registrations supported by the SBC
CurSess	—	the current call count reported by the SBC
MaxSess	—	the maximum sessions for which the SBC is licensed

Service Port Data

Service Port	—	the service path (the concatenation of realm, IP address, port number, and IP Level 4 protocol number — 17 for UDP, 6 for TCP)
App	—	the supported protocol: SIP
Handle	—	the handle for the service port
TunNdx	—	the tunnel the service port is registered for
Avail	—	current availability (yes or no) determined by the presence of heartbeats

show ccd stats

The **show ccd stats** command displays end-point statistics for the SBC members of the cluster.

```
alesmith# show ccd stats
17:10:09-54
                          -------- Period -------    ---- LifeTime ----
 SD                      Active  Rate   High Total  Total PerMax    High
bass@172.30.68.35       I285714   0.0 285714    0 285.71K 13.76K 285.71K
guinness@172.30.68.33   I285714   0.0 285714    0 285.71K 13.76K 285.71K
magichat@172.30.68.34   I285714   0.0 285714    0 285.71K 13.76K 285.71K
newcastel@172.30.68.37  I285714   0.0 285714    0 285.71K 13.76K 285.71K
samadams@172.30.68.31   I285714   0.0 285714    0 285.71K 13.76K 285.71K
sixtus@172.30.68.36     I285714   0.0 285714    0 285.71K 13.76K 285.71K
westy                   I285714   0.0 285714    0 285.71K 13.76K 285.71K

Total end-points: 153908
Total SDs      : 9
alesmith#
```

The *Period* stats provided represent an accumulation of data for the amount of time specified after the dash separator in the timestamp printed in the first line of output (in this example, the period represents 54 seconds).

The single ASCII character between the *SD* column and the *Active* column is the state of that SBC; the letter I represents *InService*.

The *Rate* column displays the transmission rate of *new* end-point associations to that particular SBC. (In the sample, no new end-points are arriving in the cluster, so all of the SBCs show a rate of 0.0.) The High field indicates the highest number of active end-point associations for the current period.

When issued with an optional *hostname* argument, the **show ccd stats** command provides a detailed report for the target hostname.

```
alesmith# show ccd stats bass
15:09:25-59
SD bass@172.30.68.33 [InService]
State                 -- Period -- -------- Lifetime --------
            Active    High   Total      Total Permax     High
Tunnels          1       1       0          2      1        1
Service Ports    2       2       0          2      1        2
end-points    53571   53571      0      53571  14399    53571
Contacts      53571   53571      0      53571  14399    53571
Sessions         0       0       0          0      0        0
Remote CPU       0       0       0          0      0        0
SD Something        ---- Lifetime ----
                       Recent          Total         PerMax
Heartbeats rcvd            30          27426             15
Heartbeats Missed           0              1              1
Tunnel Adds                 0              2              1
Tunnel Removes              0              1              1
Service Adds                0              2              1
Service Removes             0              0              0
end-point Removes           0              0              0
end-point Promotes          0          53571          13561
end-points Skipped          0              0              0
Rebalance Source            0              0              0
Rebalance Targe             0              0              0
Rebalance Request           0              0              0
```

```
Rebalance Replies           0             0             0
CPU Above Limit             0             0             0
CPU Above Threshold         0             0             0
Online Transitions          0             0             0
Offline transitions         0             0             0
Tunnel Add Fails            0             0             0
CCD Tunnel Add Fails        0             0             0
Tunnel Remove Fails         0             0             0
CCD Tunnel Remove Fails     0             0             0
Service Add Fails           0             0             0
CCD Svc Add Fails           0             0             0
Service Remove Fails        0             0             0
CCD Svc Remove Fails        0             0             0
Service Adds No Cfg         0             0             0
Bad Service Handle          0             0             0
end-point Remove Fails      0             0             0
end-point Prom Fails        0             0             0
alesmith#
```

The *Period* stats provided represent an accumulation of data for the amount of time specified after the dash separator in the timestamp printed in the first line of output (in this example, the period represents 59 seconds).

Tunnels contains the number of tunnels between the SLB and the target SBC, in this case, *bass*.

Service Ports contains the number of Service Ports advertised by the target SBC when it joined the cluster.

end-points and *Contacts* contain the number of end-point associations the SLB has assigned to the target SBC. If there is only one registering device at a given end-point, a one-to-one correlation between end-points and contacts is expected. However, if the **atom-limit-divisor** parameter has been set to a non-default value, the number of contacts exceeds the number of end-points.

Sessions contains the number of active calls.

The table below the overview data displays specific CCP message statistics.

HeartBeats rcvd contains the number of heartbeat/keepalive messages received from the target SBC. Heartbeats are sent every two seconds by the SBC.

HeartBeats Missed contains the number of scheduled heartbeat/keepalive messages not received from the target SBC.

The *Tunnel Adds* and *Tunnel Removes* counters are incremented when an SBC joins the cluster and leaves the cluster, respectively.

The *Service Adds* and *Service Removes* counters are incremented when an SBC advertises support for a service and withdraws support for a service, respectively. This generally happens only when an SBC first joins the cluster, or if the configuration on a clustered SBC is changed, saved, and activated.

The *end-point Removes* counter tracks the number of SBC-originated Cluster Control messages that request the SLB to delete a forwarding rule. Such a request can be the result of (1) a rebalance operation (when the SLB asks for the SBC to nominate candidates for rebalancing), (2) an end-point de-registration with the SBC, or (3) an end-point is power down. Generally, whenever a registration cache entry on a clustered end-point is removed by the SBC, it notifies the SLB to remove that binding.

The *end-point Promotes* counter tracks the number of promotion messages the SBC sends to the SLB to validate an untrusted forwarding rule. When the SLB first creates a forwarding rule for a new end-point, it treats it as untrusted. When the SBC receives a 200 OK for a REGISTER message from that end-point's registrar, the SBC sends a Promote Cluster Control message to the SLB. At this point, the SLB modifies the particular forwarding rule and assigns it *trusted* status. If this Promote message is not received within the time configured as the untrusted-grace-time in the *lbp-config*, the SLB deletes the untrusted entry.

end-points Skipped contains the number of end-points in its registration cache that the SBC has skipped over during a rebalance request. Skipping may be done for one of two reasons: either the most appropriate user for rebalancing was in an active phone call (and **rebalance-skip-calls** was enabled in *cluster-config*), or the **rebalance-skip-ahead** value in *cluster-config* was set to a nonzero value. In this case, when the SBC is asked to nominate users for rebalance, it will skip over any users whose registration cache entry is due to expire within the number of milliseconds set as the **rebalance-skip-ahead** value.

Rebalance Source contains the number of times the target SBC was used as a source of end-points during a rebalance operation (that is, it supplied end-points to a cluster member that was added to the cluster after itself).

Rebalance Target contains the opposite: the number of times that SBC was the recipient of end-points from other sources during a rebalance operation.

The *Rebalance Requests* and *Rebalance Replies* counters increment upon receipt of a Cluster Control message from the SLB to the SBC asking it to divest itself of end-points, and the responsive Cluster Control message from the SBC that indicates the end-points the SBC has chosen.

The *CPU Above Limit* and *CPU Above Threshold* counters increment whenever an SBC has reported a high CPU value, and has been taken out of consideration for new end-point assignments. Generally, the CPU limit and threshold are the same value (90%). However, it is possible to configure the threshold to be lower using the sip-config **option load-limit**.

SBC/Cluster Member Statistics

The SBC cluster member also provides the operator with summary statistical data for active endpoints

show sip lb-endpoints

The **show sip lb-endpoints** command displays SBC end-point stats by realm or tunneled service ports, by sip-interface since each SIP interface is uniquely identified by its realm name.

As seen below, release SC-Z 7.1.2 provides more information than was reported by prior releases.

Net-Net 4500 SC-Z7.1.2

```
westy# show sip lb-endpoints
-------------------------------------
 Realm Endpoint Stats
-------------------------------------
15:55:32-38
Service access
               ------- Period -------      ------ Lifetime -----
               Active    High   Total      Total  PerMax High
Endpoints     1000    1000       0       5020     850 1000
Atoms         1000    1000       0       5020     850 1000
                          ------- Lifetime -------
                          Recent      Total  PerMax
Refreshes                   0        999     547
Adds                        0       5020     850
Low Skips                   0          0       0
High Skips                  0          0       0
Auth Promo Tries            0        991     539
noTrust Promo Tries         0          0       0
Promo Tries                 0        991     539
Remote Deletes              0          0       0
SP Removes                  0          0       0
Expiry Deletes              0         10      10
Session Deletes             0          0       0
Session Adds                0          0       0
Move Deletes                0          0       0
Move Del No Tells           0          0       0
SvcMove Deletes             0       4010     537
SvcMove Del NoTell          0          0       0
Auth Promotes               0        991     538
Auth Deletes                0          0       0
Add Errors                  0          0       0
Delete Deny Sess            0       4230       9
Delete Deny Reg             0      11269      10
Delete Deny Purge           0          0       0
```

```
Delete Missing            0          0       0
Delete Errors             0          0       0
Update Deny Purge         0          0       0
Auth Deny Purge           0          0       0
Remote Sess Skips         0          0       0
Remote Del Fails          0          0       0
SP Remove Fails           0          0       0
Expiry Del Fails          0          0       0
Sess Del Fails            0          0       0
Sess Add Fails            0          0       0
Move Del Fails            0          0       0
Move No Tell Fails        0          0       0
SvcMove Del Fails         0          0       0
SvcMv NoTell Fails        0          0       0
Auth Promo Fails          0          0       0
Auth Del Fails            0          0       0
App Cache Dels            0          0       0
westy#
```

Net-Net 4500 SC-X6.4.0

```
westy# show sip lb-endpoints
-------------------------------------
 Realm Endpoint Stats
-------------------------------------
15:44:30-40
Service access
               ------- Period -------      ------ Lifetime -----
               Active    High   Total      Total  PerMax High
Endpoints     1000    1000       5       3190     995 1000
Atoms         1000    1000       5       3190     995 1000
                       ------- Lifetime -------
                       Recent      Total  PerMax
Refreshes               20       7029     980
Low Skips                0          0       0
High Skips               0          0       0
Auth Promo Tries        21       4010     979
Remote Deletes           0          0       0
SP Removes               0          0       0
Expiry Deletes           0       2190      10
Session Deletes          0          0       0
Session Adds             0          0       0
Move Deletes             0          0       0
Move No Tell Dels        0          0       0
Auth Promotes            0          0       0
Auth Deletes             0          0       0
Add Errors               0          0       0
```

```
Delete Deny Sess          0       4450      10
Delete Deny Reg           0      11270      10
Delete Deny Purge         0          0       0
Delete Missing            0          0       0
Delete Errors             0          0       0
Update Deny Purge         0          0       0
Auth Deny Purge           0          0       0
Remote Sess Skips         0          0       0
Remote Del Fails          0          0       0
SP Remove Fails           0          0       0
Expiry Del Fails          0          0       0
Sess Del Fails            0          0       0
Sess Add Fails            0          0       0
Move Del Fails            0          0       0
Move No Tell Fails        0          0       0
Auth Promo Fails          0          0       0
Auth Del Fails            0          0       0
App Cache Dels            0          0       0
westy#
```

Additionally, by adding the realm name as an additional argument you can get stats by service port for each service port in the specified realm.

Net-Net 4500 SC-Z7.1.2

```
westy# show sip lb-endpoints access
--------------------------------------
 ServicePort Endpoint Stats
--------------------------------------
15:57:46-53
Service Port access:192.168.168.100:5060<17>
               ------- Period -------      ------ Lifetime -----
               Active    High   Total      Total  PerMax High
Endpoints     1000    1000       0       5020     850 1000
Atoms         1000    1000       0       5020     850 1000
                          ------- Lifetime -------
                          Recent      Total  PerMax
Refreshes                   0       1999    1000
Adds                        0       5020     850
Low Skips                   0          0       0
High Skips                  0          0       0
Auth Promo Tries            0        991     539
noTrust Promo Tries         0          1       1
Promo Tries                 0        992     539
Remote Deletes              0          0       0
SP Removes                  0          0       0
Expiry Deletes              0         10      10
Session Deletes             0          0       0
```

```
Session Adds              0          0       0
Move Deletes              0          0       0
Move Del No Tells         0          0       0
SvcMove Deletes           0       4010     537
SvcMove Del NoTell        0          0       0
Auth Promotes             0        992     538
Auth Deletes              0          0       0
Add Errors                0          0       0
Delete Deny Sess          0       4230       9
Delete Deny Reg           0      11269      10
Delete Deny Purge         0          0       0
Delete Missing            0          0       0
Delete Errors             0          0       0
Update Deny Purge         0          0       0
Auth Deny Purge           0          0       0
Remote Sess Skips         0          0       0
Remote Del Fails          0          0       0
SP Remove Fails           0          0       0
Expiry Del Fails          0          0       0
Sess Del Fails            0          0       0
Sess Add Fails            0          0       0
Move Del Fails            0          0       0
Move No Tell Fails        0          0       0
SvcMove Del Fails         0          0       0
SvcMv NoTell Fails        0          0       0
Auth Promo Fails          0          0       0
Auth Del Fails            0          0       0
App Cache Dels            0          0       0
westy#
```

show sip ccp

The **show sip ccp** command displays a cluster-member-specific summary of CCP operations.

```
westy# show sip ccp
--------------------
MOO:0.4:T2
--------------------

EP Del                Recent      Total  PerMax
===================   ======  =========  ======
Ops Sent                   0          1       1
Op Replies Recvd           0          1       1
```

```
                      -------- Received ------- ---------- Sent --------
Status Code           Recent      Total  PerMax Recent      Total  PerMax
-----------          ------  ---------  ------   ------ ---------  ------
200 OK                     0          1       1 0          0       0

EP Promo              Recent      Total  PerMax
===================   ======  =========  ======
Ops Sent                   0        992     538
Op Replies Recvd           0        992     538

                      -------- Received ------- ---------- Sent --------
Status Code           Recent      Total  PerMax Recent      Total  PerMax
-----------          ------  ---------  ------   ------ ---------  ------
200 OK                     0        992     538 0          0       0

Metrics               Recent      Total  PerMax
===================   ======  =========  ======
Ops Sent                  25        207      15
Op Replies Recvd          25        207      15

                      -------- Received ------- ---------- Sent --------
Status Code           Recent      Total  PerMax Recent      Total  PerMax
-----------          ------  ---------  ------   ------ ---------  ------
200 OK                    25        207      15 0          0       0

Stop Down             Recent      Total  PerMax
===================   ======  =========  ======
Ops Sent                   0          2       2
Op Replies Recvd           0          2       2

                      -------- Received ------- ---------- Sent --------
Status Code           Recent      Total  PerMax Recent      Total  PerMax
-----------          ------  ---------  ------   ------ ---------  ------
200 OK                     0          2       2 0          0       0
westy#
```

Subscriber-Aware Load Balancer SNMP Reference

Overview

This chapter provides an overview of SNMP support for Subscriber-Aware Load Balancer (SLB) features.

Enterprise Traps

The following table identifies the SLB proprietary traps supported by the SLB.

apSLBEndpointCapacityThresholdTrap	Generated when the number of endpoints on the SLB exceeds the configured threshold.
apSLBEndpointCapacityThresholdClearTrap	Generated when the number of endpoints on the SLB falls below the configured threshold.
apSLBUntrustedEndpointCapacityThresholdTrap	Generated when the number of untrusted endpoints on the SLB exceeds the configured threshold.
apSLBUntrustedEndpointCapacityThresholdClearTrap	Generated when the number of untrusted endpoints on the SLB falls below the configured threshold.

License MIB (ap-license.mib)

SNMP GET Query Name	Object Identifier Name: Number	Description
Object Identifier Name: apLicenseEntry (1.3.6.1.4.1.9148.3.5.1.1.1)		
apLicenseSLBEndpointCap	apLicenseEntry: 1.3.6.1.4.1.9148.3.5.1.1.1.23	SLB endpoint capacity

Subscriber-Aware Load Balancer MIB (ap-slb.mib)

SNMP GET Query Name	Object Identifier Name: Number	Description
Object Identifier Name: apSLBMIBObjects (1.3.6.1.4.1.9148.3.11.1)		
Object Identifier Name: apSLBMIBGeneralObjects (1.3.6.1.4.1.9148.3.11.1.1)		
apSLBStatsEndpointsCurrent	apSLBMIBGeneralObjects: 1.3.6.1.4.1.9148.3.11.1.1.1	Number of endpoints currently on the SLB.
apSLBStatsEndpointsDenied	apSLBMIBGeneralObjects: 1.3.6.1.4.1.9148.3.11.1.1.2	Number of endpoints denied by the SLB because the system has reached the maximum endpoint capacity.
apSLBEndpointCapacity	apSLBMIBGeneralObjects: 1.3.6.1.4.1.9148.3.11.1.1.3	Maximum number of endpoints allowed on the SLB. This value is based on the installed SLB license(s).
apSLBEndpointCapacityUpperThresh	apSLBMIBGeneralObjects: 1.3.6.1.4.1.9148.3.11.1.1.4	Percentage of endpoints relative to maximum threshold capacity.
apSLBEndpointCapacityLowerThresh	apSLBMIBGeneralObjects: 1.3.6.1.4.1.9148.3.11.1.1.5	Percentage of endpoints relative to minimum threshold capacity.
apSLBStatsUntrustedEndpointsCurrent	apSLBMIBGeneralObjects: 1.3.6.1.4.1.9148.3.11.1.1.6	Number of untrusted endpoints currently on the SLB.
apSLBStatsTrustedEndpointsCurrent	apSLBMIBGeneralObjects: 1.3.6.1.4.1.9148.3.11.1.1.7	Number of trusted endpoints currently on the SLB.
apSLBStatsUntrustedEndpointsDenied	apSLBMIBGeneralObjects: 1.3.6.1.4.1.9148.3.11.1.1.8	The number of untrusted endpoints denied by the SLB due to the total number of untrusted endpoints exceeding the configured maximum threshold.
apSLBStatsUntrustedEndpointsAgedOut	apSLBMIBGeneralObjects: 1.3.6.1.4.1.9148.3.11.1.1.9	The number of untrusted endpoints aged out of the system because they were not authenticated within the configured grace period.
apSLBUntrustedEndpointCapacity	apSLBMIBGeneralObjects: 1.3.6.1.4.1.9148.3.11.1.1.10	Maximum number of untrusted endpoints allowed on the SLB. This value is a configured percentage of the maximum endpoint capacity of the system.
apSLBUntrustedEndpointCapacityUpperThresh	apSLBMIBGeneralObjects: 1.3.6.1.4.1.9148.3.11.1.1.11	Percentage of untrusted endpoint maximum threshold capacity in use.
apSLBUntrustedEndpointCapacityLowerThresh	apSLBMIBGeneralObjects: 1.3.6.1.4.1.9148.3.11.1.1.12	Percentage of untrusted enpoint minumum threshold capacity percentage.