

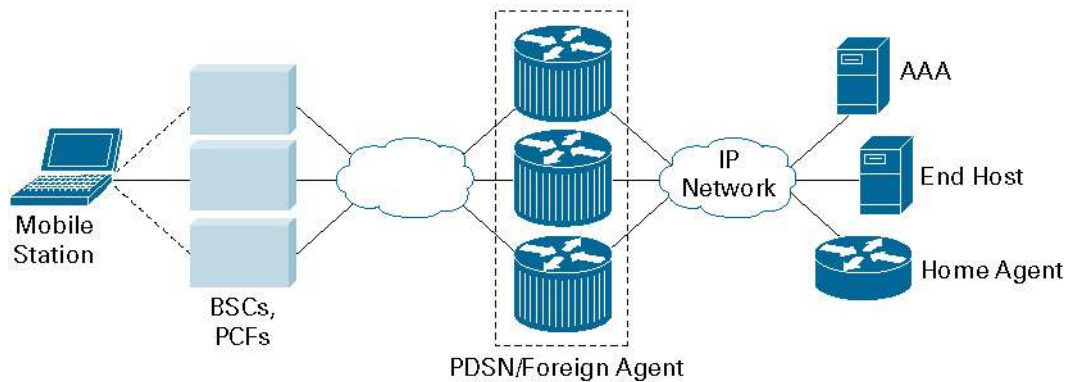
CISCO PDSN CLUSTERING

Packet data serving nodes (PDSNs) can be configured in a group to form a logical PDSN cluster as illustrated in Figure 1. Allowing the flexibility of adding PDSNs transparently, the clustering concept assures optimal performance as well as providing a scalable approach to growing the network. As demand increases, additional PDSNs can be added to a cluster without disrupting any service and assuring load is distributed evenly across all PDSNs. Advantages of clustering include the following:

- Enhanced PDSN availability
- PDSN redundancy and resiliency
- PDSN load balancing
- Minimized inter-PDSN handoffs, including for simple IP user

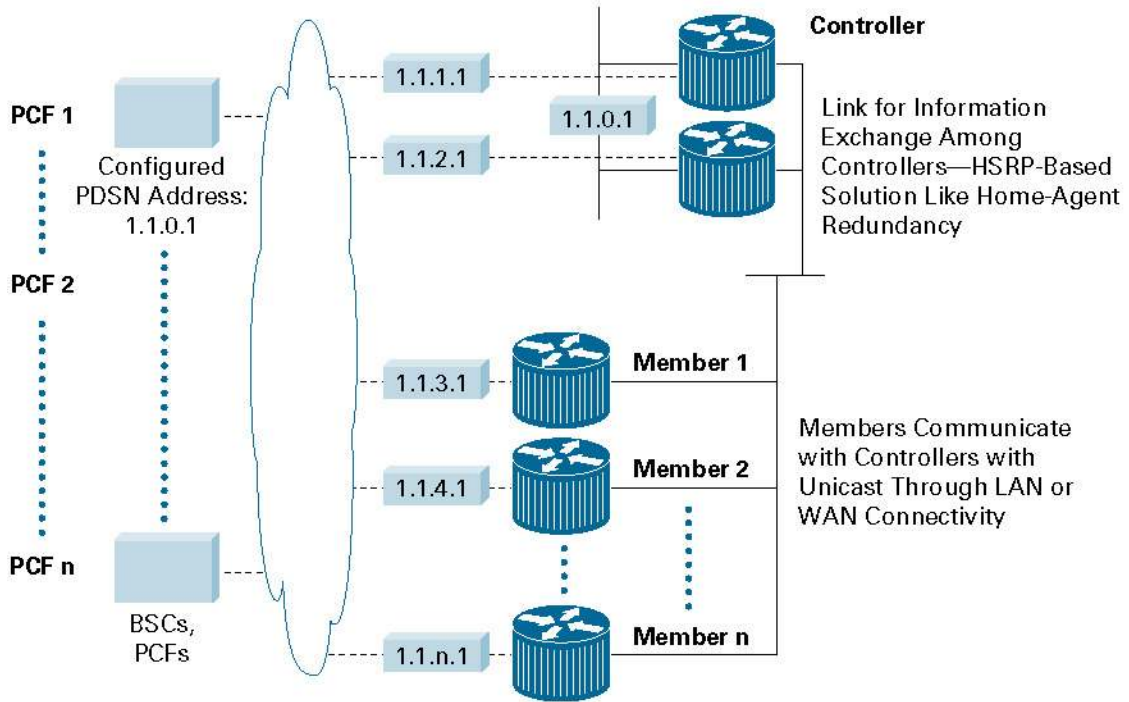
Figure 1

Shows Grouping of PDSNs to Form a Logical Cluster



Cisco® PDSN supports a clustering technique named “PDSN Cluster Controller Member Architecture.” In this architecture the concept of “controller” and “member” functions is introduced, resulting in a mixed hierarchical and peer-to-peer model. Controllers are responsible for determining the optimum-member PDSN for a given session and directing the session to that PDSN. Members are responsible for reporting a measure of their ability to host sessions in addition to the primary purpose of functioning as a PDSN. A representation of the PDSN cluster in the controller member architecture is shown in Figure 2.

Figure 2
Shows PDSN Cluster with Controller and Members

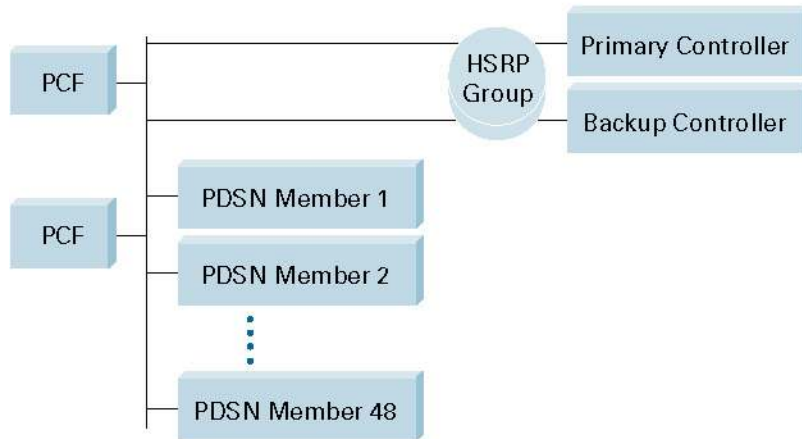


CISCO PDSN CLUSTER CONTROLLER MEMBER ARCHITECTURE

The Cisco PDSN cluster controller member architecture was introduced in PDSN Release 1.2 to satisfy the clustering requirements resulting from the improved scalability performance of the release. The architecture embraces the notions of controllers and members, each with distinct functional responsibilities.

A controller is responsible for maintaining load and session (such as an A10 connection) information for each member in a cluster and performs member selection for load balancing or inter-PDSN handoff avoidance. A controller is responsible for identification of the operational state of each member and detection of a member failure. Two controllers, one active and one standby, can be configured as a Hot Standby Router Protocol (HSRP) group for redundancy. The role of a controller is just that: to control a group of PDSN members. Figure 3 illustrates a logical representation of the cluster.

Figure 3
Cisco PDSN Cluster



A member is responsible for notifying the active controller about its load and session information and providing PDSN control and bearer-plane functions.

The initial A11 registration request for a session is sent from the packet control function (PCF) to the active cluster controller. The IP address used is that of the HSRP group containing the cluster controllers. Based on load information maintained in the cluster controller, a cluster member is chosen. An A11 registration reject is sent to the PCF with a reject reason of 0x88, 136 decimal, and the selected cluster member's IP address in the home-agent information element. The PCF is then expected to send a new registration request for that session to the indicated cluster member.

MECHANISM AND MESSAGES

Between PCF and Cluster

A cluster can have up to 48 PDSN cluster members. Only 1 controller is required, but most installations will use 2 controllers (1 active and 1 standby for 1 + 1 redundancy) and up to 48 members.

The controller group is configured with HSRP. The active controller in the controller group provides throughput to process A11 registration requests for up to 960,000 Point-to-Point Protocol (PPP) sessions. The active controller, which knows the load and session information of each member in the cluster, receives and processes all the A11 registration request messages sent by a PCF. The standby controller has no role in that process. Upon receipt of an A11 registration request, the active controller selects a member and, by default, responds with an A11 registration reply message to the sending PCF with the code 0x88 and the IP address of the selected member. The cluster is transparent to a standards-compliant PCF. The PCF finishes the rest of session setup A11 message exchange with the selected member PDSN and tunnels the A10 data between the PCF and the member. All subsequent A11 signaling for this A10 connection happens directly between the PCF and the member.

The controller logical address is entered into the list in a PCF that it uses for requesting new tunnels. When a request is made to a PDSN cluster controller for a new tunnel, the PDSN controller accepts the request or denies it and suggests another PDSN (that is, a member within the cluster) to serve the request. The controller chooses the member to serve the request to balance the subscriber call load across PDSN members and to minimize inter-PDSN handoffs. If the cluster is fully loaded (all PDSNs within the cluster are serving the maximum number of allowed calls), then the PDSN controller rejects the request indicating that the cluster is at full capacity.

A PCF initiates a request to establish a RP tunnel under the following conditions:

- A new data session is initiated by the subscriber
- The PCF is the target of an inter-PCF handoff

When a PCF establishes a data session, it sends an A11 registration request to one of the PDSNs on its list. The PDSN(s) on the PCF list are PDSN cluster controllers. How a particular (controller) PDSN is selected from the list of (controller) PDSNs is a PCF function and depends on the PCF implementation or vendor.

The load-indicator value represents a percentage of the member usage. Although the load indicator on every member of one cluster is maintained approximately uniform, that load level may vary from cluster to cluster.

Between Controller Logical Entity and Members

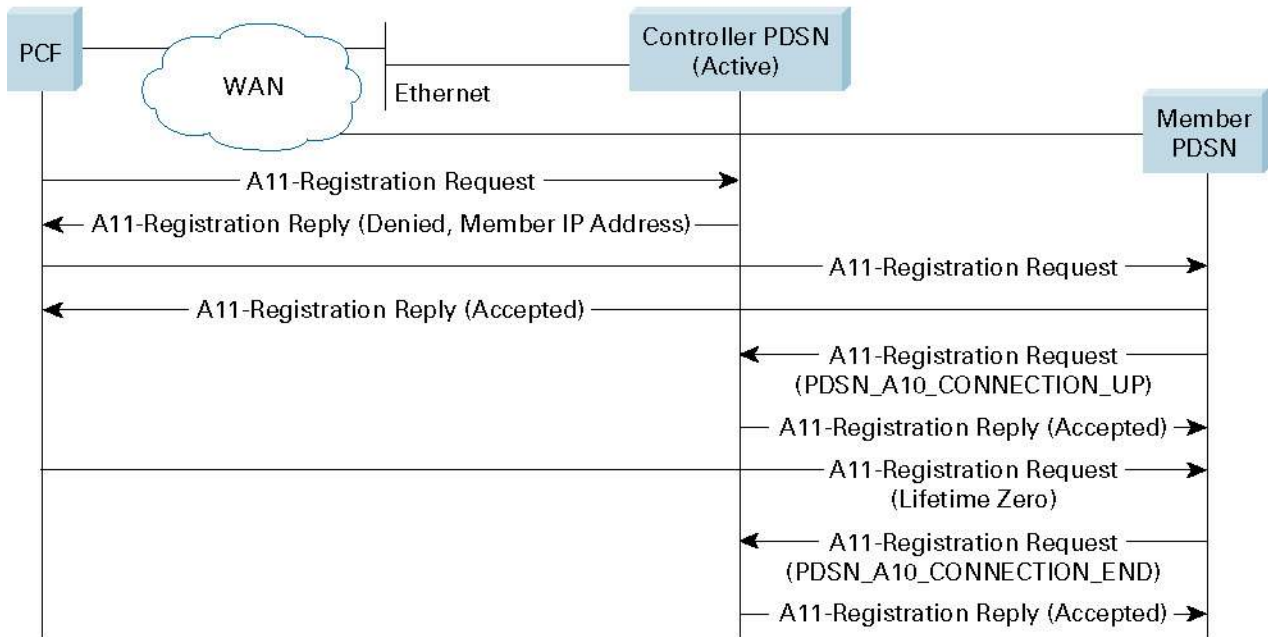
Messaging between a controller and its members is designed to be reliable and minimal but sufficient.

Upon the end of configuration, the members send an A11 registration request message to the configured IP address of a controller, and continue until the controller replies. In this way the member joins the cluster.

Normally, members send messages related to session origination and termination and the controller replies; if the messages are frequent enough there is no need to send any other message. Only at periods of low control traffic will an active controller possibly need to seek its members to verify if they are online. A “seek” is an A11 registration request. Seeks are reliable. They are transmitted a configurable number of times if no response comes from a member. Eventually, if there is no member response, the member is presumed offline.

The controller calculates load information of members using the information received from the members as a percentage of their usage. The data is contained in a vendor-specific extension inside an A11 registration request or reply message. A member times out controller replies or seeks. If a timer expires the member concludes the controller has failed and starts again to seek its controller by sending an A11 registration request message and retransmits it until a reply comes back. Figure 4 illustrates the generic establishment and teardown call flows.

Figure 4
Shows Generic Call-Flow Establishment and Teardown



1. A11 registration requests go from the PCF to the active controller PDSN. The active controller PDSN looks up a session table using the incoming mobile-station identifier (MSID) as the index. If an entry is present for the MSID controller, the PDSN forwards the registration request to the member PDSN that hosts the session for the MSID. If the session entry is not present in the controller PDSN session table, it chooses a member PDSN based on the configured selection algorithm.
2. The active controller PDSN replies to the PCF with an A11 registration reply (with error code 0x88 as defined in TIA/EIA/IS2001) suggesting the member PDSN IP address.
3. The PCF sends an A11 registration request to the suggested member PDSN when the session comes up on the member PDSN, a session-up message is sent to the controller PDSN from the member for that session (MSID).
4. On receipt of the session-up message from the member PDSN, the controller PDSN creates an entry in the session table for that MSID to establish an MSID-member PDSN association on the controller.
5. If the PCF sends an A11 registration request with lifetime zero, the session is brought down on the member PDSN and a session-down message is sent to the controller from the member PDSN.
6. On receipt of the session-down message from the member PDSN, the controller PDSN flushes the entry for the MSID from the session table.

REDUNDANCY

Cluster redundancy is based on the premise that there could be at most a single PDSN failure at any given time. Two cluster controllers are configured as an HSRP group; one controller is active, the other standby. Cluster-member redundancy is an N + 1 scheme. If a cluster member

fails, the established sessions are lost but the overall group capacity is such that the sessions can be reestablished with the other group members. Loss of a session can be avoided by taking advantage of another feature called session redundancy.

An additional aspect that enhances redundancy is that cluster members no longer have to be network neighbors.

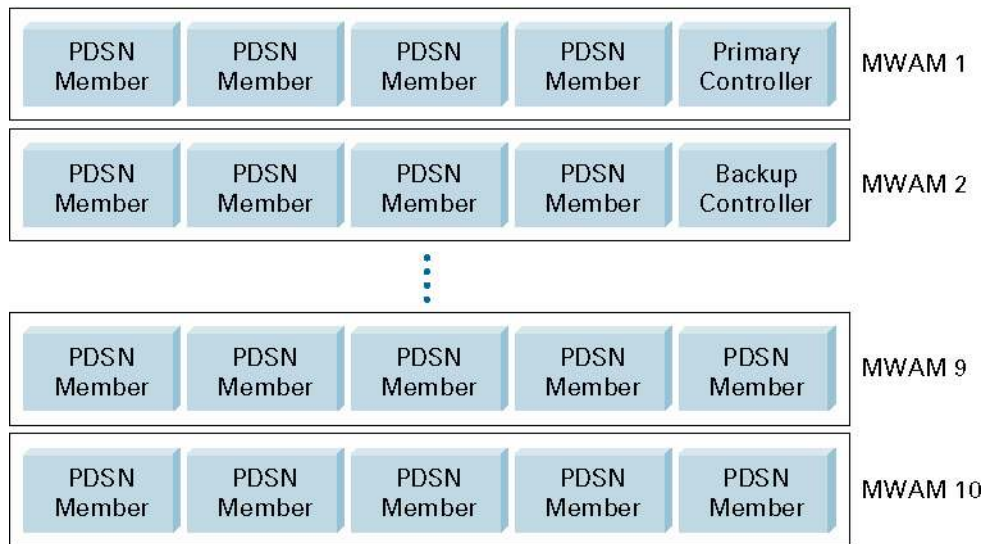
Controllers exchange information on an Ethernet link.

The session and member records are synchronized between the active and standby controllers as needed. Because all controllers of a cluster maintain session and load information for all the members of that cluster, failure of a controller does not result in the loss of any session or load information.

The link on which the controller-member information exchange takes place is a unicast interface over which the members address messages to the HSRP group address of the controllers. The PDSN members in a cluster do not need to be network neighbors. This means that members can be attached anywhere in the IP network, thereby enhancing the geographical redundancy characteristics of the cluster.

Adding an additional controller to a cluster is simplified by autoconfiguration of the controller in the cluster. This is possible by configuring the additional controller for HSRP. The newly added controller then synchronizes automatically with the active controller. Similarly, when a new member is added to the cluster, autoconfiguration for the member occurs in all cluster controllers. Figure 5 illustrates a full cluster with 2 PDSN controllers and 48 PDSN members.

Figure 5
Shows Full PDSN Cluster (2 Controllers and 10 Members Across 10 MWAM Blades)



The application blade can be within a single chassis or across two or more chassis.

PDSN CLUSTER-MEMBER SELECTION

Selection of a cluster member by the controller is based on the load factor. Load factor is the ratio of the number of sessions supported to the total session capacity of the PDSN. The controller attempts to assign sessions to members so as to achieve a uniform load factor for all members.

If an A11 registration request is received indicating handoff, a member that is already servicing the session is selected. This applies independently of inclusion of the mobility event indicator (MEI) in the registration request.

LOAD BALANCING

A controller maintains load information for all members in the cluster in order to perform the PDSN cluster-member selection. This information is transferred from the members to the controller:

- At periodic intervals
- When an accumulation of session data results in the member informing the controller in which case the periodic timer is restarted
- When the controller requests it for the members

Members send the load-factor information to the controller in an A11 registration request containing a vendor-specific extension.

A controller requests load information by performing a seek operation. A seek is achieved by sending an A11 registration request with a different vendor-specific extension to the cluster member. A seek is transmitted a configurable number of times, and no response after that implies member failure.

The seek rate is based on a timer, the “seek timer.” Receipt of any session-related message from the member results in a restart of the seek timer, hence delaying transmission of the seek. In this way, if a member is periodically sending session-related messages to the controller, a seek is not required. All session-related messages contain the load-factor information. Figure 6 shows a cluster with eight registered members with their corresponding load and remaining time before next polling (seek) with as well the logical IP address of the controller (IP address of the HSRP group to which belong the controllers).

Figure 6

Sample Display of Eight Members in a Cluster with the Corresponding Load, Seek Information, and Logical IP Address of the Cluster

sh cdma pdsn cluster controller member load				
Secs until (past) seek	Seq seeks no reply	Member Ipv4 Addr	State	Load
3	0	10.10.95.1	ready	4
2	0	10.10.93.1	ready	4
8	0	10.10.92.1	ready	4
6	0	10.10.94.1	ready	4
5	0	10.10.72.1	ready	4
3	0	10.10.74.1	ready	4
9	0	10.10.75.1	ready	4
9	0	10.10.73.1	ready	4
Controller Ipv4 Addr 10.121.68.98				



FINE-TUNING

In Cisco PDSN Release 1.2, when the session comes up, the member sends a session-up message to the controller. The controller receives a session-down message when the member's session is closed. This results in two message exchanges between the controller and the member for every session opened.

In Cisco PDSN Release 2.0, a new feature is added to reduce the number of message exchanges and hence increase the number of members supported per cluster. To reduce the processing overhead members, send one bulk-update packet to the controller for every configured periodic update time interval with multiple pairs of session-up or session-down messages. The packet contains concatenated multiple MSIDs with one session-up or session-down flag. The controller processes these bulk-update packets and sends a bulk-update acknowledgment packet to the members.

**Corporate Headquarters**

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters

Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters

Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the **Cisco Website** at www.cisco.com/go/offices.

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica
Croatia • Cyprus • Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR
Hungary • India • Indonesia • Ireland • Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico
The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal • Puerto Rico • Romania • Russia
Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden • Switzerland • Taiwan
Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

Copyright © 2004 Cisco Systems, Inc. All rights reserved. Cisco, Cisco Systems, and the Cisco Systems logo are registered trademarks or trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0406R)
He/LW7112 09/04