

Using the Cisco Catalyst 3750 Series Switch as a Gigabit **Linux** Cluster Interconnect

The advent and growth of the server cluster market has been dramatic, especially with the stability, portability, and cost savings that the Linux operating system provides. Although Linux clustering offers a cost-effective method for organizations to perform large computations, its effectiveness depends highly on the type of interconnect used between the servers in the cluster. This white paper describes the specific advantages of using the Cisco Catalyst® 3750 Series Switch as a Gigabit interconnect in Linux server clusters.

Overview of the Cisco Catalyst 3750 Series Switch as an Interconnect

The new Cisco Catalyst 3750 Series switches improve LAN operating efficiency by combining industry-leading ease of use and the highest resiliency available for stackable switches. This new product series represents the next generation in desktop switches, and features Cisco StackWise™ technology, a 32-Gbps stack interconnect that allows customers to build a unified, highly resilient switching system—one switch at a time.

Server Clusters

A Linux server cluster is a collection of inexpensive computers running Linux that work together to perform large-scale computing jobs. Each individual Linux computer in the cluster is known as a *node*; sometimes they are also defined as individual processors or as individual servers. By providing high-speed communication (such as Gigabit Ethernet) between these nodes, it is possible to separate a large job or several jobs and assign sections to each individual node.

Due to their reliability and computational power, mainframes and supercomputers were used to perform these large-scale jobs prior to the advent of server clusters. However, this reliability and computational power came at a high cost. In addition to the significant initial costs of these systems, adding capacity to a mainframe or supercomputer was complicated and expensive; these large systems required customized physical installations.

Linux server clusters now offer the same advantages with none of the drawbacks of mainframes and supercomputers. With today's technology, it is possible to build a powerful cluster, including the display and keyboard, that is entirely contained in a standard, 19-inch, data-center equipment rack. The U.S. government currently operates the largest known cluster (more than 5000 nodes), and the size and computing power of future server clusters are limitless.



There are typically two types of Linux server clusters. The first type is a high-availability cluster that is used to offer redundant data on multiple nodes. These clusters provide automated failover to minimize or eliminate outages due to system failure. The second type is a parallel-computing or high-performance cluster. This type of cluster is intended for applications that involve numerous calculations that are not interdependent. There are two different versions of the parallel-computing cluster type, one focused on capability and one focused on throughput.

- Capability clusters address huge problems that need to scale to multiple CPUs. It is more efficient and economical to solve some of these problems in a cluster environment than on a single large system.
- Throughput clusters run multiple jobs (either different applications or multiple instances of the same application) in a batch-like environment.

The media used for interconnection heavily influence the effectiveness of the cluster.

The Cluster Interconnect

Server clusters have used several different media types for interconnection. They differ in the amount of bandwidth they offer, the latency they introduce, and in their total cost. The most popular modes include:

1. *Gigabyte System Network (GSN)*—This method of interconnect offers the highest bandwidth (6.4 Gbps) and lowest latency. GSN is used in clusters that require the most bandwidth and lowest latency, such as capability clusters for scientific simulation. It is expensive and cost-prohibitive for most conventional cluster uses.
2. *Myrinet*—Myrinet offers lower bandwidth (2 Gbps) and higher latency than GSN. Although still expensive, it is less so than GSN and therefore more applicable to conventional clusters. Until the advent of Fast Ethernet as an interconnect, it was the only choice for conventional cluster applications.
3. *Gigabit Ethernet/Fast Ethernet*—Fast Ethernet is a relatively inexpensive means of interconnecting. Because of the cost savings achieved by using Fast Ethernet, it has grown in popularity as an interconnect protocol. Clusters with relatively low-bandwidth requirements that do not require low latency, including many throughput and high-availability clusters, have found Fast Ethernet to be the best and most economical interconnect choice.

Gigabit Ethernet is the newest method of cluster interconnects. It offers good bandwidth (1 Gbps) and tolerable latency. The cost for Gigabit Ethernet has dropped to nearly the same cost as Fast Ethernet—far more affordable than GSN or Myrinet—making Gigabit Ethernet a viable option. With the 32-Gbps Cisco StackWise interconnect, the Cisco Catalyst 3750 Series Switch becomes an even more attractive interconnect option.

Cisco Catalyst 3750 Series as a Cluster Interconnect

Compared to GSN and Myrinet systems, Cisco Catalyst 3750 Series switches offer several advantages as a deployment option for cluster interconnects (Table 1).



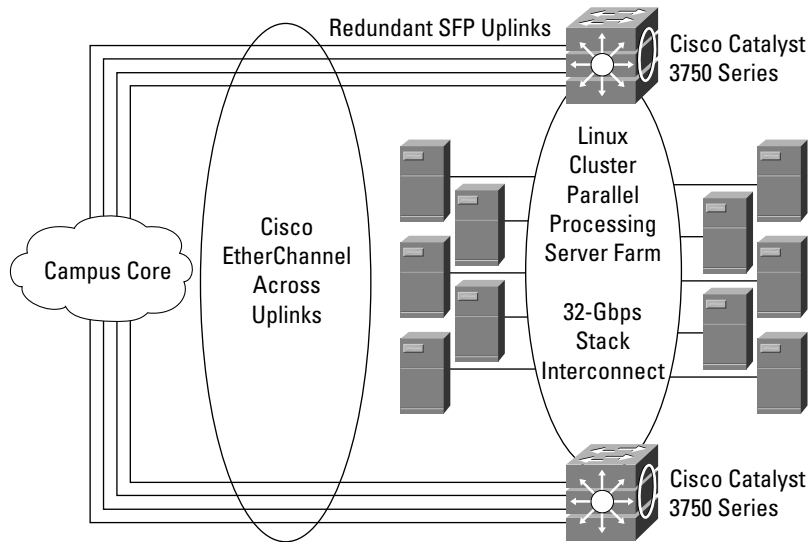
Table 1 Features and Benefits of the Cisco Catalyst 3750 Series in Linux Cluster Environments

Cisco Catalyst 3750 Series Features	Customer Benefit
<p>High-speed 32-Gbps stack interconnect</p> <ul style="list-style-type: none"> • Uses high-performance Cisco StackWise technology 	<ul style="list-style-type: none"> • Allows data from several Gigabit Ethernet links to servers to flow through the stack concurrently • Allows for greater densities of servers in a cluster • Larger switch fabric (by several multiples across the stack than conventional stackable switches)
<p>Line-rate performance</p> <ul style="list-style-type: none"> • Leads to low latency through the switch • Line rate performance, even with quality of service (QoS) and security activated 	<ul style="list-style-type: none"> • Satisfies one of the basic requirements for a cluster interconnect—low latency • QoS and security features help ensure rapid handling of packets, while adding useful functions
<p>QoS and policing to prioritize traffic</p>	<ul style="list-style-type: none"> • Offers preferential treatment for certain data streams, if required • Segments traffic streams into different paths for processing, if required
<p>Cisco EtherChannel[®] technology across stack</p>	<ul style="list-style-type: none"> • Provides protection against device failure • Provides bandwidth augmentation and redundancy for links across the entire stack • Provides redundant uplinks to the network core (Figure 1)
<p>Lower cost</p>	<ul style="list-style-type: none"> • Cisco Catalyst 3750 Series switches, with Fast Ethernet or Gigabit Ethernet technology, offer a less-expensive solution for cluster interconnect than GSN and Myrinet
<p>Architectural advantages</p>	<ul style="list-style-type: none"> • The 48-port gigabit platform in the Cisco Catalyst 3750 Series offers 4 Small Form-Factor Pluggable (SFP) uplinks, allowing dual redundant uplinks as an additional availability advantage. • The various lengths of stack cable available, ranging from 0.5 to 3 meters (approximately 1.5 to 10 feet), allow the stack to be extended across multiple racks for multiple stack aggregation

Figure 1 illustrates a possible implementation scenario, with a stack of Cisco Catalyst 3750 Series switches as the interconnect mechanism for a cluster of Linux servers.

Figure 1

Enterprise Server Clustering Using the Cisco Catalyst 3750 Series



Using Cisco Catalyst 3750 Series switches, customers can take full advantage of the cost savings of Gigabit Ethernet as a server cluster interconnect, without compromising bandwidth or latency.



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.
Capital Tower
168 Robinson Road
#22-01 to #29-01
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 Web site at www.cisco.com/go/offices**

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia
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

All contents are Copyright © 1992–2003 Cisco Systems, Inc. All rights reserved. Cisco, Cisco Systems, the Cisco Systems logo, Catalyst, Cisco IOS, EtherChannel, and StackWise are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

All other trademarks mentioned in this document or Web site 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.
(0303R) WH/LW4394 0403