ipXpress

by Phobos Corp.

IP**X**PRESS™
Web Traffic Load Balancing Appliance

Load balance multiple servers in a cluster.

- **Balances** web traffic among all servers.
- **Easily connects** to any Ethernet network.
- **Quick** set up and remote configuration.
- **Choose** from Six different load balancing algorithms.
- **Hot** standby failover port for web site uptime.

Web sites that consist of multiple servers or server farms need web traffic to be balanced among all the servers for optimal performance of the site. Phobos' ipXpress™ is a four-port load balancer that connects quickly to any Ethernet or Fast Ethernet network, automatically switching traffic to the most available server based on 6 different load balancing options. Small and sleek with highly visible status and activity lights, the ipXpress is the ideal product for colocation cages and rack rooms alike.

Once installed, setup is accomplished through remote configuration software that can be run from any Linux or Windows NT computer on the network. A failover port lets you connect to a second ipXpress via serial cable, so you'll always have a hot standby if one device goes down.

The ipXpress has a powerful on-board processor running Phobos' own MaxOS operating system to perform all switching and load balancing functions. The ipXpress continually monitors the health of each server to ensure availability and peak performance. Six different algorithms are available to determine which server should receive requests: Round Robin, Least Connections, Weighted Least Connections, Fastest Response Time, Adaptive and Fixed.

For high-availability of the web site, the ipXpress can be configured to support failover from primary to secondary servers. If a server goes down, the ipXpress will route traffic to redundant servers until the failed device comes back up.

sslXpress

SSL**X**PRESS™

eCommerce Accelerator Appliance

Perform secure transactions at lightning speed.

- **Accelerates** eCommerce sites up to 50 times.
- **Offloads** secure transactions, freeing up servers.
- **Simple** installation on any Ethernet network.
- **Guaranteed** security for all SSL transactions.
- **Compatible** with load balancing appliances.

According to Internet analysts The Industry Standard (January 2000), Andersen Consulting found that 88 percent of online buyers abandoned their electronic shopping carts at some time during the 1999 holiday season. The heavy processing burden of Secure Sockets Layer (SSL) transactions used for encryption of credit card information can overwhelm secure servers, bringing web sites to a crawl and frustrating would-be buyers.

Phobos sslXpress™ improves the performance of eCommerce sites by helping them transact secure business and deliver sensitive information quickly, confidentially and without errors. The sslXpress performs all key management and encryption of Secure Sockets layer (SSL) transactions, off-loading this burden from web servers. The result is a tremendous boost for busy eCommerce sites, improving performance up to 50 times.

This small, two-port appliance is ideal for eCommerce sites in colocation cages. Highly visible status lights and a convenient reset switch facilitate monitoring and troubleshooting of the device. Remote configuration and maintenance of the sslXpress can be performed from any PC running Windows NT or Sun Solaris operating systems. A failover port lets you connect to a second sslXpress via serial cable, so you'll always have a hot standby if one device goes down.

The sslXpress works in conjunction with IP (Internet Protocol) traffic load balancers, improving the overall performance of any eCommerce web site. By installing an sslXpress in front of the load balancer, all resource-intensive data encryption and decryption functions are performed before traffic is distributed to web servers in the cluster. This enables the load balancer to perform advanced functions such as URL parsing and cookie-based load balancing.

Part#	Product	Price
IP400	ibXpress	$ 3,495
SL200	sslXpress	$ 4,995

What is the PHOBOS ipXpress?

The ipXpress is a network appliance that balances web traffic efficiently among all the servers in a web server farm, ensuring the high availability and reliability of your web site.

We are a small company but are quickly growing our web presence. Why do we need a load-balancing product?

Many of our customers have grown from single server to multi-server web sites. The ipXpress lets your web site grow at your own pace. As demand for your products and services increase, you can easily add additional web servers using the ipXpress. Another point to consider is the reliability and resilience gained from a load-balancing solution such as the ipXpress. For example you can connect two redundant web-servers to the ipXpress and allow one to act as a backup in the case the other fails or goes down.

Is the ipXpress hardware-based or software-based?

The ipXpress is a hardware device. It is a small four-port switch that connects to an Ethernet or Fast Ethernet network between the router and the web servers.

We have Windows NT and SUN Solaris web servers. What OS does the ipXpress support?

Because the ipXpress is a separate network appliance, it supports web servers running any operating system, using a variety of Internet applications and services over TCP/IP.

Our web site is a combination of static pages, audio clips, and applications that access backend databases. What type of Internet applications and servers does the ipXpress work with?

The ipXpress supports Internet applications such as web browsers (HTTP), file transfer (FTP), mail (SMTP) and other IP (Internet Protocol) traffic.

How does the ipXpress perform dynamic load balancing?

The ipXpress uses TCP port and IP address information to direct connection requests based upon user-configured preferences. Six dynamic load-balancing algorithms or methods are available to meet your specifications.

- ***Round Robin: distributes connections evenly across the servers that it manages. Each time a new connection is requested, the ipXpress passes the connection to the next server in line. The Round Robin algorithm treats all servers within the group as equals, regardless of the number of connections or response time. Round Robin server selection is the fastest of all the algorithms.***
- ***Least Connections: directs network connections to the server with the least number of open connections. In maintaining the same number of connections to all servers, those capable of processing connections the fastest will get the most connections. Least Connections is most effective when the cluster of***

servers is similar in performance. This algorithm determines the next available server at the application level by its ability to process and terminate connections.

- *Weighted Least Connections: a refinement of the Least Connections option. This algorithm allows you to assign a performance weight to each server. Servers with a higher weight value will receive a greater percentage of connections when the number of open connections is equal. For example, a server with a weight of 3 will receive three times as many connections as a server with a weight of 1. By assigning weights to the physical servers in a group, the administrator can guarantee that faster servers get more traffic.*
- *Fastest Response Time: passes a new connection to a server based on the fastest measured response time of all currently active servers. Response time is determined by measuring the time that elapses between sending each packet to the server and receiving each packet from the server. Web server performance does not usually follow a linear progression of response time to number of connections. Web servers seem to respond flatly and then, at a certain load level, a dramatic increase in the response time occurs. The Fastest Response Time algorithm can eliminate these problems by ignoring the number of connections and relying on server speed.*
- *Adaptive: directs network connections to the server with the least number of open and pending connections. Pending connections are those that have not completed the full TCP handshake. The Adaptive algorithm performs best as the number of client threads per second increases.*
- *Fixed: the ipXpress uses the source IP address of each incoming request to decide which server receives the request. With a large number of requests from a mixture of addresses, this algorithm works well. The Fixed Mode algorithm performs better with a mixed source of requests rather than a large number of requests coming through the same gateway.*

How many virtual IP addresses and real IP addresses can the ipXpress support?

You can set up an unlimited number of virtual IP addresses. The ipXpress supports up to 8,192 physical servers in any number of server groups.

Does the ipXpress monitor server health?

Yes. The ipXpress polls each physical server to determine its state. These polls follow user-configured preferences. If a physical server is determined to be "down", the ipXpress does not direct traffic to it.

How does the ipXpress ensure that services are available?

The ipXpress uses several methods to make sure that physical servers are available:

1. ***The ipXpress "pings" physical servers. The time interval between pings is user-configurable as well as the number of retries before the physical server is considered unavailable.***

2. ***The ipXpress monitors TCP connections. You can specify which ports and which physical servers to monitor. Application server status records are logged into the host system.***
3. ***Active Content Verification: The ipXpress monitors all active connections, checking all return codes. Error codes are entered into the host system log. In the event a server is deemed inactive, busy, or not responding; the ipXpress takes the following actions:***
    - ***server taken out of active group***
    - ***server monitored to determine when active again***
    - ***server placed back in active group***
4. ***The ipXpress redirects TCP connections when a physical server is no longer responding to a client. As with many ipXpress preferences, this communication specification is use configurable.***
5. ***The ipXpress allows an administrator to flag a physical server for maintenance. No new connections are sent to that server, and after all sessions have ceased, the administator is notified that the server may be taken down for maintenance.***
6. ***The ipXpress will reset all connections going to an unavailable port. No client connections will be lost.***

What is Layer 4 switching?

Layer 4 switching is the process used by a device to make a switching decision, forwarding packets based on session (transport layer of Layer 4) information.

Can server load balancing, Layer 2 switching, and web cache redirection be performed simultaneously?

Yes.

Does the ipXpress support persistent HTTP or "sticky" connections for any protocol?

Yes. This is a user-configurable option available on a per-port basis.

Can other types of traffic besides HTTP be redirected?

Yes. The ipXpress can redirect any type of TCP traffic based upon user-configurable preferences.

What if I need to bring down a server for maintenance? Will all of the connections it is handling be lost immediately?

No. The ipXpress configuration software allows the administrator to flag any server as being unavailable. The ipXpress will stop directing connections to that server and notify the administrator when all sessions are complete and the physical server may be shut down.

ipXpress SPECIFICATIONS

Model Number
IP400

Standards
IEEE 802.3 10Base-T Ethernet
IEEE 802.3u 100Base-TX/FX Fast Ethernet

Protocols
CSMA/CD Ethernet

Data Transfer Rates

Ethernet	Fast Ethernet
10 Mbps (half duplex)	100 Mbps (half duplex)
20 Mbps (full duplex)	200 Mbps (full duplex)

Topology
Star

Network Cabling
Cable type: Category 5 UTP (100m)
Connector type: RJ-45

Number of Ports
10/100 TX 4 ports

Media Interface Exchange
All ports are configured in crossover switch arrangement

Memory
32 Megabytes RAM, 2 Megabytes Flash ROM

Configuration Software OS Support
Linux Red Hat versions 5.2 and 6.0 and Windows NT Server 4.0

Load Balancing Algorithms

- Round Robin
- Least Connections
- Weighted Least Connections
- Fastest Response Time

Adaptive
Fixed

Warranty
2 Years

Compliance
Radiated Emissions: FCC Class A, CISPR-22, VCCI and CE Compliant.

PHYSICAL AND ENVIRONMENTAL

Product Size
Dimensions: Length 9.25 inches, Height 2.3 inches, Width 8.3 inches

Weight
2 lbs.

Port Descriptions
Network Ports: 4 10/100BaseTX Ethernet ports

Failover Port: Molex Micro-Fit Serial port

Reset Switch: Push to reset hardware. Configuration data maintained.

Environmental Operating Range
Temperature: 0^{o} to 50^{o} C (32^{o} to122^{o} F)

Humidity: 10 to 85% non-condensing

Altitude: Up to 3,048 meters (10,000 ft)

Power Requirements
Operating voltage: 5 VDC 5 Watts

Power Supply: Uses Phobos external power supply only (P/N 47-0004)