



jetNEXUS Accelerator in Operation



Table of Contents

1.	INTRODUCTION	1
2.	HOW WILL JETNEXUS ACCELERATOR BENEFIT YOU?	2
3.	HOW JETNEXUS ACCELERATOR WORKS	3
4.	JETNEXUS ACCELERATION IN OPERATION	5
5.	ADMINISTRATION.....	6
6.	ADVANCED FEATURES	7

Notices & Disclaimers

The information contained within this document is the property of Metal Minds Ltd (Metal Minds), is issued in confidence and must not be discussed with or given or communicated to any third party without the prior permission of Metal Minds. No part of this document may be reproduced, stored in a retrieval system or transmitted in any form or by any means, whether electronic, mechanical, photocopying, recording or otherwise without the express written permission of Metal Minds. Any pricing information contained in this document is in the nature of a preliminary estimate supplied for budgetary purposes only and shall not be construed as a formal offer open for acceptance. This document does not form or constitute part of a contractual document, nor does its submission imply acceptance of any commercial terms.

Any advice given or statements or recommendations made shall not in any circumstances constitute or be deemed to constitute a warranty by Metal Minds as to the accuracy of such advice, statements or recommendations. Metal Minds shall not be liable for any loss, expense, damage or claim arising out of the advice given or not given or statements made or omitted to be made in connection with this document.

Copyright © 2005 by Metal Minds Ltd

THIS DOCUMENT CONTAINS 10 PAGES INCLUDING TITLE PAGE

1. Introduction

jetNEXUS Accelerator is a web acceleration appliance that builds on Metal Minds' highly successful jetNEXUS acceleration engine for IIS web servers. It is designed to bring content acceleration to any web server or HTTP-based service and to deliver the benefits of significantly reduced bandwidth usage along with dramatically improved page download times.

jetNEXUS Accelerator applies maximum acceleration by ensuring the right content is accelerated for the requesting browser. Decompression is handled by the client browser (nearly every browser in use today has standardised decompression built-in), so no client plug-in is required and deployment is transparent to site visitors.

jetNEXUS Accelerator can be used by e-commerce sites, corporate data centres, extranets, web hosts, service providers, as well as applications dependent on HTTP services. It is platform-independent and applies high performance content compression to web services such as iPlanet/SunOne, WebLogic, WebSphere, Apache and Tomcat (among others) along with their supporting database and application service environments.

Most HTTP transmissions (e.g. ASP, JSP, XML, CSS, JavaScript, DHTML, HTML, etc.) can be compressed by up to 90% with a typical bandwidth reduction of 50% or more.

If you have any questions, or would like to trial jetNEXUS Accelerator please do not hesitate to contact us using the details provided.

2. How will jetNEXUS Accelerator Benefit You?

MS Internet Explorer and Netscape Navigator have had, since version 4, the ability to accept some level of accelerated content WITHOUT requiring a plug-in or other modification to the browser. Web servers on the other hand have themselves provided little or no direct support for this ability. jetNEXUS Accelerator provides the means to exploit this feature, and in doing so significantly improves the customer experience as well as providing considerable network benefits.

Once installation has taken place, jetNEXUS will provide your business with an instant reduction in bandwidth requirement (directly proportional with the level of acceleration achieved). Your customers will experience improved page delivery times, so that both you, and your customers, will experience the benefits of jetNEXUS Acceleration.

As jetNEXUS Accelerator is a separate hardware device, there is no adverse effect on server resources (CPU and Memory Usage). jetNEXUS Accelerator integrates easily into any infrastructure offering the instant benefits of compression without the requirement for user plug-ins, or changes to content.

jetNEXUS Accelerator has a scalable, modular design which allows web sites to evolve over time. Advanced features offered by jetNEXUS Accelerator include:

- File specific acceleration techniques
- Active grid failover
- SSL connection termination
- Load balancing
- HTTP acceleration performance parameters

These features will be described later in this document.

The primary management interface for jetNEXUS Accelerator is a web-based console. The web console requires no special management software other than a standard browser. The interface is intuitive and easy-to-use. A basic configuration can be up and running within minutes. At the same time the web console is powerful enough to allow advanced configuration tailored to the requirements of a complex site. The web console supports all of the advanced features provided by jetNEXUS Accelerator.

3. How jetNEXUS Accelerator Works

jetNEXUS Accelerator is delivered as a dedicated HTTP 1.1 compliant appliance. Incorporating support for multiple Virtual IP addresses and an automated configurable failover feature jetNEXUS Accelerator has a scalable, modular design, which is applicable to different web site architectures and network topologies. jetNEXUS Accelerator is compatible and complimentary to static edge caching technologies and load balancers, such as Akamai, Digital Island, Alteon, Cisco and F5 Networks.

The fundamental process is as follows:



Fig 3.1 Client Request through jetNEXUS Accelerator

The user enters a URL or clicks on a link. This creates a connection followed by a request from the browser to the web site. As part of the HTTP protocol, the browser identifies in its request headers whether it is capable of receiving compressed content as well as the methods of compression it understands.

On receiving the client browser's request, jetNEXUS Accelerator stores the request's HTTP header information and passes the request on to the site's web service. Once the request has been forwarded, jetNEXUS Accelerator uses this data to determine the browser's ability to receive compressed content and, if appropriate, prepares to compress the outbound content.

The site's web service then processes the HTTP request and generates the appropriate content, usually in conjunction with dynamic data sourced from back-end systems. jetNEXUS Accelerator is not sensitive to any of these downstream activities and is compatible with web services such as iPlanet/SunOne, WebLogic, WebSphere, Apache and Tomcat (among others) along with their supporting database and application service environments.




Fig 3.2 Server Response through jetNEXUS Accelerator

The web page requested by the client browser is then forwarded to jetNEXUS Accelerator for processing. It examines the response headers and determines whether the output is of a compressible type. If it isn't (e.g. GIF, JPEG, etc.), then jetNEXUS Accelerator stays out of the way. However, if the output is of a compressible type AND jetNEXUS Accelerator knows from the request that the browser can accept that content in a compressed form, jetNEXUS Accelerator modifies the response. The content of the response is compressed and the response headers are altered as required.

On completion, the page is forwarded to the client browser for decompression and rendering. From the end user's point of view this whole process is transparent, apart from the fact that the page arrives much, much faster.

jetNEXUS Accelerator is optimised to communicate with all common browsers and has in-built knowledge of which browsers can handle accelerated content and for which page elements. For example, jetNEXUS will accelerate CSS style sheets when communicating with Internet Explorer, but not with early Netscape browsers. JetNEXUS Accelerator is also aware

of any acceleration-relevant software defects associated with particular browser releases as well as their remedy. Where the browser is unable to support compressed output, jetNEXUS Accelerator simply stays out of the way.

4. jetNEXUS Acceleration in Operation

For a site being accessed by Internet users, jetNEXUS Accelerator would typically be located just inside the site's router/firewall and in front of its web service as the following diagram illustrates:




Fig 4.1 Implementation of jetNEXUS Accelerator in a typical web environment

As part of a larger web service environment, jetNEXUS Accelerator would be integrated in the following manner using multiple Virtual Acceleration Engines on the same physical hardware:




Fig 4.2 Implementation of jetNEXUS Accelerator using multiple Virtual Accelerator Engines

Further examples of how jetNEXUS Accelerator can be used as part of a redundant web site architecture are given in the Active Grid Failover discussion within the Advanced Features section of this document.

5. Administration

Installation

For easy installation onto your network, jetNEXUS Accelerator comes with the Windows based jetNEXUS Discovery program. jetNEXUS Discovery makes any jetNEXUS Accelerators on your local network instantly visible, allows easy configuration of the administration IP address and gives convenient access to the web console.

As delivered, jetNEXUS Accelerator is configured with sensible defaults that allow a basic configuration to be up and running within minutes.

Web Console Features

The web console supports configuration of all the major operational features of jetNEXUS Accelerator, including:

- Acceleration Engine listening configuration
- Active grid failover
- SSL Termination
- Load balancing
- HTTP acceleration performance parameters
- Exclusion of specific files from the compression process

In addition, the web console includes the following management capabilities:

- At-a-glance system status monitoring
- Detailed performance statistics
- Configuration backup and restore
- Software version upgrade
- System security configuration
- SNMP setup
- W3C and system log access

If required the web console can operate over an SSL interface.

W3C Logging

A common problem with any piece of software or hardware located between the client browser and the web server is that client identification can be lost because the HTTP request appears to have been made by an internal site device. jetNEXUS Accelerator offers alternative solutions to this problem.

Firstly jetNEXUS Accelerator can be configured to include an X-Forwarded-For header, containing the client address, in the HTTP request that is sent on to the web server.

Alternatively, the W3C logging activity can be offloaded from the web server to the jetNEXUS accelerator. The W3C log file can be stored locally on the jetNEXUS accelerator or can be compressed and copied to a remote file server.

SNMP

jetNEXUS Accelerator includes support for SNMP V1/V2/V3 and allows basic system status monitoring via standard MIBs.

6. Advanced Features

Multiple Acceleration Engines

jetNEXUS Accelerator will support multiple Virtual Acceleration Engines on the same physical hardware, with each Virtual Acceleration Engine listening on a different IP address and TCP port combination.




Fig 6.1 Implementation of jetNEXUS Accelerator using multiple virtual accelerator engines

In the above example a single jetNEXUS Accelerator is located in front of multiple web servers. Within the jetNEXUS Accelerator a separate virtual acceleration engine will map to each of the web servers, forwarding on HTTP requests to that web server and accelerating responses from that web server. When another web server is introduced, configuring a new virtual compression engine to map to the new web server is a straightforward task and the new virtual acceleration engine will be operational in less than 60 seconds. This approach allows Accelerator resources to be dedicated to each web server, maximising throughput and minimising response times.

Failover

Redundancy and fault tolerance are necessities for any commercial web site. The jetNEXUS Accelerator Failover feature allows active and standby Virtual Accelerator Engines to co-exist on the same physical hardware. This means that in a redundant configuration equipment is not left idle, and can contribute to the overall performance of the site. This is a far more efficient use of resources than a passive failover configuration, as described below.




Fig 6.2 Passive failover configuration with jetNEXUS Accelerator

Fig 6.2 shows a passive failover configuration. A backup jetNEXUS Accelerator is configured with an identical set of Virtual Acceleration Engines. If the active jetNEXUS Accelerator fails, the backup jetNEXUS Accelerator provides a drop in replacement.

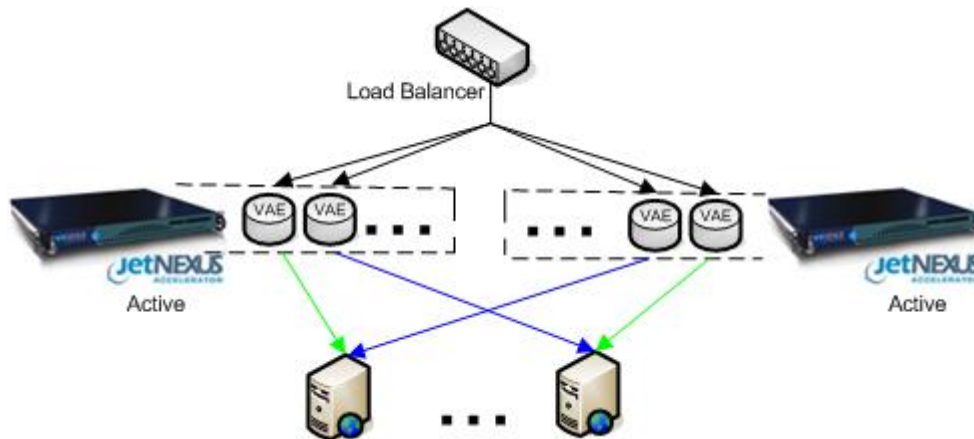


Fig 6.3 Active grid failover configuration with jetNEXUS Accelerator

Fig 6.3 shows how active grid failover provides a more cost effective and higher performance alternative. In this configuration both jetNEXUS Accelerators are active simultaneously. For each virtual acceleration engine, one jetNEXUS Accelerator will be identified as the primary. This means that under normal operation the virtual acceleration engine will be active on that jetNEXUS Accelerator. The jetNEXUS Accelerators will each monitor the health of the other jetNEXUS Accelerators. If one of the jetNEXUS Accelerators detects a failure in any of the virtual acceleration engines for which they are not the primary, the standby virtual acceleration engines on the secondary Accelerator will be brought online. This solution is automatic and requires no manual intervention when failures are detected. When the primary virtual acceleration engine is able to restore service, the secondary virtual acceleration engine will be returned to a standby state and the primary virtual acceleration engine will be come active again.

SSL Termination

There are two options for termination of SSL connections.

1. SSL connections can be terminated by the web server. The advantage of this option is that no extra equipment is required. A disadvantage of this approach is that a busy web server will be impacted by the CPU overhead required by the SSL activity. Also, it will not be possible to compress the encrypted server responses because the content will now be a string of ever changing characters, so the overall compression benefit is reduced.
2. jetNEXUS Accelerator can terminate SSL connections to reduce the processor overhead of SSL activity on the web server. This also allows the compression activity to occur before SSL encryption, maximising the amount of data that can be compressed. This solution is also cost effective because no extra equipment is required.

Configuration of the optional SSL Termination feature within jetNEXUS Accelerator is fully supported by the graphical user interface. This includes the ability to upload certificates provided by a Certificate Authority and to generate self signed certificates.

Load Balancing

jetNEXUS Accelerator can provide a cost-effective acceleration and load balancing solution for the small to medium sized company. In a load-balanced configuration, a Virtual Accelerator Engine will map to a group of web servers. When a HTTP request is received from a client, the load balancing policy will select a content server to forward the request onto and the response will be accelerated and returned to the client as normal. The policy used by jetNEXUS Accelerator load balancing can be selected from options that currently include round robin, least connections and sticky sessions based on client IP address. The load-balancing feature is fully supported by the web console.