Features	Functions	Benefits
Server Load Balancing	Layer 4 - 7 load balancing	Operating at the application protocol level increases functionality, delivering features such as URL switching, cookie based persistence and request scanning. This content-aware load balancing strategy enables application performance optimisation and resilience across the application layer. Deploying the jetNEXUS ETM in a high availability N+M cluster eliminates any single point of failure, guarantees uninterrupted service delivery and provides immense scaling for performance.
	Slow-start server re-introduction	
	Customisable server too busy page	
	Automated per-channel fast failover	
	High availability deployment	
	Connection draining	
Load Balancing Methods	Round robin	Traffic is distributed across client's server pool via multiple load balancing strategies for maximum flexibility in deployment. Application cookie sticky method uses existing application cookies for session persistence.
	Least number connections	
	IP sticky	
	Cookie sticky	
	Fastest response time	
	Application cookie sticky	
Session Persistence	Pre-defined and customisable methods	Persistent load balancing methods for session based applications including ecommerce. Automatic session detection dynamically sets up cluster-aware persistence.
	HTTP and SSL specific persistence modes	
Content Caching	GUI based cache control rule base	jetNEXUS ETM cache stores common responses to web requests on behalf of the web server. Content caching thereby reduces server load and bandwidth consumption and accelerates applications. The ETM caching feature is easy to configure with multiple caching profiles.
	Multiple caching profiles	
	Configurable expiry	
	Cache hit reporting	
	Persistent cache	
Server Health Checking	Ping server health check	The ETM continually monitors the health and status of servers to detect and route around problem servers – from simple ping and TCP tests to full web GETs. Automatic failover to enable clients to deliver a seamless and fully fault-tolerant service delivery platform. Recovered servers are automatically and gradually reintroduced.
	TCP connect server health check	
	Simple HTTP server health check	
	Full HTTP server health check	
	Full fault reporting for failed transactions	
Service Protection	Web worm, DoS and DDos protection	Limiting of client concurrent connections and connection rates protects servers against Denial of

	Real-time attack monitoring and threat analysis	
	Per-service access controls	Services and Distributed Denial of Service attacks.
	Policy-based request filtering	
	Connection limiting	Server Concurrency Limiting prevents server overload and provides optimal request distribution.
	Real-time attack monitoring and threat analysis	Heuristics protect against malformed URL attacks, buffer over-run attempts and invalid HTTP requests
	Configurable attack logging	IP black-and white-lists for basic access control policies.
SSL Termination	High performance SSL offload	
	Native SSL decryption	SSL Offload removes the overhead of CPU intensive encryption from the web servers to minimise the performance impact on the server pool.
	Over 8,000 tps	
	SSL re-encryption to back-ends	This optimizes server performance and enhances the end user experience. The ETM has high performance software SSL stack optimized for x86 and SPARC processors. SSL decryption and back-end encryption delivers end-to-end security.
	Centralised certificate management	
	Ability to create self signed cert	
Management	Active-Active failover configuration	
	Connection proxy capability	
	Secure web based interface	Secure, resilient web-based GUI with wizards to simplify common tasks. Tools for easy management and flexible configuration and changes.
	SNMP (V1,2 &3)	
	Online configuration file back up	SOAP API for remote and automated management.
	Software updatable via GUI	
	Command line interface CLI	Simple, fast deployment with jetNEXUS Discovery Tool for automatic detection.
	Email alerting and real time stats	
	DNS resolution	Integration with enterprise authentication services for administration.
	Dedicated management interface	
	Optional client IP x forwarding	Catalogs for easy storage and re-use of common configuration.
	W3C web transaction logging & file offload	
	SOAP API	Performance monitoring with customizable real-time analysis and traffic visualization and trending for intelligent management.

Content Acceleration	Dynamic HTTP compression	Compressing content before it is sent to clients helps to reduce the burden on the networking infrastructure. On-the-fly content compression is applied to any compressible content type by the jetNEXUS ETM. Rule based compression improves performance and delivers bandwidth savings.
	Content exclusions – by rule	
	Configurable compression	
	Personal firewall compression	
	Browser rule base compression	
	Point to point compression-decompression	
	Streaming compression	
Connection Management	Connection management	Optimises the performance of TCP/IP and manages stress on web servers. Key to high performance load balancing.
	Connection pooling	
	Connection capping/limiting	
RuleBuilder	GUI-based traffic management rules	Create bespoke traffic management rules for easy and intelligent service management. Rules are stored in the catalog for re-use and easy deployment to multiple virtual servers.
	Content-based traffic routing	
	URL, HTTP header and HTTP cookie rewriting	
	Rules are stored in the Catalog for re-use	
Java Extensions	Customisable with JAVA extension	Develop traffic management policies using Java or any JVM-based language Allows for use of any java class libraries, e.g database access, document watermarking etc.
TrafficScript	Inspect, manipulate and route transactions	A feature of the ETM, TrafficScript is a powerful routing engine. **Use cases include:** • Analyse and rewrite client requests and server responses to optimise transactions, detect protocol and security anomalies, work around server errors and to perform response assembly. • Base traffic decisions on origin, destination, content type or any part of the transaction data. • Prioritize and manage traffic based on context, e.g. transaction value, client location, user activity, service performance. • Apply session persistence and routing based on any parameter or value in the request.
	Support for TCP and UDP protocols	
	Protocol specific functions	
	Unlimited content inspection depth	
	Native processing of XML data using Xpath, XSLT and Schema and DTD validation	

Optional Feature	Functions	Benefits
Service level Monitoring	Create user-defined service levels	With Service Level Monitoring, clients can define and create acceptable levels of performance and monitor these service levels on a real time basis. User defined alerting and custom remedial actions
	Real-time monitoring and alerting	

	Monitor using GUI, SOAP &SNMP	can be set, should performance fall outside of service level limits. Clients can also apply traffic management policies selectively based on service performance e.g to prioritise key transactions and limit others.
	Apply policies based on service performance	
	Configure auto remedial actions	
Bandwidth Management	Apply per-service bandwidth limits	By applying bandwidth limits on a per-connection, per group of connections or per service basis, this feature helps to maintain a high level of service for key users. This feature is especially useful in ecommerce environments whereby the client may wish to restrict the impact of busy peak periods.
	Active, cluster-aware, real-time bandwidth management	
	Restrict impact of high service demand	
	Apply bandwidth classes selectively	
Real-Time Analytics	Record and visualise full transaction headers	Analytics is an optional feature of the ETM, enabling users to understand end-user transactions and their effect on applications in real-time. Real-time analytics enables operational staff to drill down into recent transactions and gain detailed timing information to identify issues and rapidly fix problems.
	Identify anomalies and faults	
	Full visibility of client-server interactions	
Rate Shaping	Rate shape client requests within Rules	Individual users may dominate the use of a service, to the detriment of other users of the service. A back-end application infrastructure may have limited scalability, being easily overwhelmed when too many requests are given to it. You may wish to restrict the rate at which certain activities can occur, such as sending an email, or logging in to a service, as part of a wider security policy.
	Specify limits on a wide range of events	
	impose per-second and per-minute rates	
Multi-Site Manager	Centrally manage a cluster of ETM appliances over one or more sites	Multi-Site Manager extends the administration capability of jetNEXUS Traffic Manager by allowing the large multi-site cluster to span multiple different locations. The extended administration interface can control which configuration is deployed in each location, and provides a global view of performance, health and status from the fault-tolerant, fully distributed management interface.
	Replicate service configuration to all in the cluster	
	Ability to set local-specific configuration	
Application Auto-scaling	Monitors response time from pool and provides scale-up/scale-down thresholds	Monitors the performance of a service running on a supported platform. When the performance falls outside the desired service level, ETM can then initiate an auto-scaling action, requesting that the platform deploys additional instances of the service. ETM will automatically load balance traffic to the new instances as soon as they are available.
	Configurable stability timer	
	Create cloud credentials for cloud API	
Web Application Firewall	Licensed component of ETM	Enterprise-level Web Application Firewall that provides attack detection and protection for the latest generation of mission-critical web applications. It enables centralized security monitoring, reporting and alerting and provides custom protection for your Web applications and infrastructure against external attacks.
	Detection and Protection modes	
	Baseline and Application-based Protection	
Global Load Balancing	Datacenter load balancing	Where inter-site resilience is required, Global Load Balancing connects your users to the most suitable datacenter or cloud based on various factors such as geographic proximity, datacenter performance and datacenter availability. With optional "no-fail-back" for controlled transfer of services from one location to another.
	Geographic proximity	
	Adaptive load balancing	