

Yunke YK-LTM / YK-DNS Application Delivery Controller

YK-LTM r4800

Product Introduction

Beijing Digital China Yunke Information Technology Co., Ltd.
http://www.yunke-china.com

Product Overview



The ever-changing market environment and business demands are gradually leading to an increase in application size. As the “middleware” between enterprise applications and end users, the application delivery process requires significant enhancements in both performance and functionality.

The Yunke R-series application delivery controller solutions offer a smoother user experience with faster, more reliable, and flexible characteristics. By bridging the gap between traditional and modern infrastructure through a restructured “priority design API” model, it meets the needs of both legacy and emerging applications. In terms of performance parameters, the Yunke R-series achieves unprecedented growth. With a fully automated architecture and the highest levels of reliability, security, and access control, enterprises can shorten time to market, integrate infrastructure, and protect critical data at scale, while reducing total cost of ownership (TCO) and ensuring future-proof application infrastructure.

This document describes the YK-LTM r4800 configured as a Local Traffic Manager (LTM) with the DNS / Global Server Load Balancing (GSLB) module licensed. Web Application Firewall (WAF) capability is not included in this configuration and is available separately as the YK-AWAF product.

Features

The next-generation Yunke R-series application delivery controller provides the most outstanding high performance and scalability to date, achieving the goal of reducing application delivery quantity and lowering data center costs even in the most demanding application delivery scenarios.

Included features:

Local Traffic Manager · DNS / GSLB · SSL Offload

Optional add-on features:

Web Application Firewall · Remote Access · Network Firewall · IP Intelligence

Modern Architecture Platform Design

The Yunke R-series application delivery controller features a virtualized architecture based on Kubernetes, allowing enterprises to launch multiple Yunke application delivery tenants within the same device to meet topological needs. Multi-tenant software versions do not need to be the same, and the deployment,

upgrade, and repair processes can be carried out independently, with internal configurations, data, management access, and traffic direction all existing separately.

Automated Rapid Deployment Components

The R-series platform is designed with a fully automated platform that through a simple yet powerful declarative interface helps reduce manual operations, quickly load, support overall configuration, and deploy Yunke application services, shortening deployment time from weeks to hours.

High-Speed SSL Offloading

For critically important enterprise data, the R-series platform provides SSL offloading capabilities that are twice as effective compared to previous products, supporting forward encryption extensions, enhanced ECC elliptic curve algorithms, and full-stack protection from layers 3 to 7. By offloading costly SSL processing, it accelerates key exchange and bulk encryption speeds, speeding up SSL/TLS applications. Deliver the SSL capacity required to protect critical data—including enhanced offload of Galois Counter Mode (GCM), Elliptical curve cryptography (ECC), Camellia, and Rivest Shamir Adleman (RSA) processing to hardware—enabling forward secrecy scaling.

FPGA Assistance to Enhance CPU Performance

Utilizing widely adopted Field Programmable Gate Arrays (FPGAs), it works closely with the Yunke operating system and platform layer to enhance CPU resource utilization, providing exceptional performance and higher scalability.

Integrated DNS and Global Server Load Balancing

The same platform delivers authoritative DNS services and global server load balancing. A BIND 9 based name server, full DNSSEC signing and validation, DNS caching and DNS express zone transfer operate alongside the local traffic management functions, removing the need for separate DNS appliances and allowing GSLB decisions to use live health and load data gathered from the local traffic manager.

Flexible Migration Expansion

Through Yunke migration tools, enterprises can smoothly and frictionlessly migrate existing operations from older versions of Yunke application delivery software (both chassis and virtual versions) to the new R-series platform devices. The tool supports a graphical interface, migrating all layer 4 to 7 configurations, and supports cross-platform migration.

Visual Interface Operations

A web-based graphical interface provides enterprises with a more intuitive operation method, completing most configuration management tasks and enabling manual monitoring of current Yunke application delivery software performance.

Customizable Programmable Controllers

To leverage cloud and software-defined architecture advantages and expand application services on demand, Yunke provides fine-grained traffic observation and control. Programmable controllers allow users to easily customize code, quickly respond to code errors and security vulnerabilities in real-time, and support new protocols, facilitating integration into orchestration systems. It supports extensions to JavaScript, reducing deployment costs and accelerating deployment speed.

Value Proposition

Excellent Performance Parameters

Up to 50 Gbps / 40 Gbps of layer 4 / layer 7 throughput, 3.5M layer 4 HTTP requests per second, 1.8M layer 7 requests per second and 38M layer 4 concurrent connections, achieving industry-leading encryption performance of up to 45K TPS (RSA 2k), 20K TPS (ECDHE-ECDSA P-256) and 20K TPS (ECDHE P-256-RSA 2k) through SSL FPGA hardware acceleration, together with 30 Gbps of hardware compression offloading.

Lower Total Cost of Ownership (TCO)

Offloading application and DNS services onto high-performance Yunke application delivery devices optimizes network architecture, enhances server efficiency, and reduces overall ownership costs and infrastructure expenses. A single pair of appliances provides both local traffic management and global DNS load balancing.

Efficient and Flexible Deployment

Enterprises do not need to build pure software infrastructure; they can directly utilize and implement CI/CD toolsets, declarative APIs, and telemetries. With the “priority design API” architecture, Yunke provides a fully automated system that offers the agility and reliability enterprises need. Efficiency features include 80 Plus Platinum certified power supplies as well as front-panel touchscreen LCD management, remote boot and multi-boot support, and USB support.

Reduced Operational Difficulty

The declarative API significantly lowers operational complexity, making it easier to reduce the knowledge reserve requirements for technical personnel, relieving pressure on technical departments, and minimizing operational steps to reduce the risk of errors caused by repetitive actions.

Technical Parameters

Features	Detail Description
Main Load Balancing Features	
Server Load Balancing	Comprehensive layer 4/7 switching capabilities, supporting customizable application-layer health checks, session persistence based on IP address, cookies, etc., and health checks for associated applications, and application connection state management.
Multi-Link Load Balancing	Supports intelligent routing for multiple links simultaneously, including inbound/outbound bidirectional multi-link load balancing and rich algorithms.
Source-In Source-Out Load Balancing	Under network IP overlap, data can be returned without specifying routing, based on different source interfaces.
System Redundancy	Supports Active-Active and Active-Standby redundancy modes, providing dedicated hardware serial heartbeat lines and network-level redundancy detection.
Session Mirroring	Supports mirroring of connection and persistence information to the peer device in a device service clustering configuration, for both active/active and active/passive deployments, so that established sessions survive a failover. SSL connection and session mirroring is supported.
Programmable Traffic Management	Management interface provides customizable traffic control methods based on TCL programming language.

Features	Detail Description
API Interface	Supports REST API mode for device operations and queries, including accounts/passwords, interfaces, routing, software upgrades, and configuration files.
Programmable Control Based on Node.js Environment	The device comes with built-in programmable control features that support a Node.js environment and JavaScript syntax.
Multi-tenant	Supports multi-tenants with no fewer than 4 tenants on r4800; the tenant is allocated with independent hardware resources, and resources are isolated between tenants.
Server Load Balancing Algorithms	Supports load balancing algorithms such as Round Robin, Least Connections, Ratio, Fastest, Observed and Predictive.
Server Priority Groups	Servers providing the same service can be assigned different priorities. When the high-priority group falls below the required load, servers from the lower-priority group are activated.
Session Persistence	Supports session persistence based on source IP, cookie, destination address, SSL session ID, SIP, universal (iRule based) and host. Source IP address persistence of LDNS is supported for DNS/GSLB deployments.
Server Health Check	Supports multiple active health check methods, including ICMP, TCP, TCP half-connection, UDP, HTTP, HTTPS, FTP, SNMP, WMI, MSSQL, DNS, NTP, etc. Supports passive health check methods, monitoring real-time data between clients and servers. If application access failures are detected, user requests can be redirected to another server while marking the faulty server as unavailable. Supports complex check methods that simulate actual user access, allowing for verification of the entire transaction process from login to application access and logout. Health monitoring covers devices, applications and content.
Multipath Connection Multiplexing	Combines multiple requests from a single user or multiple users into one connection sent to the server, reducing pressure on application servers and improving response. TCP connection multiplexing is supported.
Memory Cache	Utilizes memory to cache frequently accessed web static content, reducing pressure on application servers and enhancing user response speed. HTTP RAM cache is supported.
NAT Technology	Supports NAT and PAT technologies, allowing a single IP to connect to different destination addresses using different NAT IPs. A SNAT can map many original addresses to a smaller set of translation addresses.
Link Aggregation	Supports bundling of multiple link bandwidths and supports the LACP protocol.
IPv6 Support	Supports IPv6/v4 dual stack. To set up Yunke to function as an IPv4-to-IPv6 gateway, you create a load balancing pool consisting of members that represent IPv6 nodes, then create a virtual server that load balances traffic to those pool members.
DNS and Global Server Load Balancing Features	
IPv4 and IPv6 GSLB	Supports IPv4 and IPv6 global server load balancing, including resolution of AAAA (IPv6 address) records as well as A records, and operation of the GSLB service over both IPv4 and IPv6 transport.
Load Balancing of DNS Servers	Supports load balancing of back-end DNS servers, distributing incoming DNS queries across a pool of name servers with health monitoring, so that failed name servers are removed from service automatically.
Built-in BIND 9 DNS Server	Built in with a BIND 9 DNS server, allowing the device to act as a standalone authoritative name server. Supports record types including A, AAAA, CNAME, DNAME, HINFO, MX, NS, TXT, SOA and SRV.

Features	Detail Description
DNSSEC	Supports the DNSSEC security protocol, including on-the-fly (real-time) DNSSEC signing of records returned from the authoritative name server, so that responses generated dynamically by GSLB are signed as they are served. Supports DNSSEC key management with zone-signing and key-signing keys and automatic key rollover.
DNSSEC Secure Delegation	Supports DNSSEC zone delegation to a secure (signed) child zone, publishing the DS records required to maintain an unbroken chain of trust from parent to child zone.
EDNS Client Subnet	Supports association with the Extended DNS (EDNS0) client subnet option, so that GSLB topology decisions can be based on the subnet of the originating client rather than the address of the local DNS resolver.
Zone Transfer	Supports zone transfer from an authoritative DNS server, operating as a secondary name server that copies zone information from the master DNS server and responds to client queries, protecting the back-end master server from direct query load. Both AXFR and IXFR transfers are supported.
DNS Caching	Supports DNS caching for authoritative DNS records, and DNS caching for forwarding (resolver) DNS queries. Cached responses are served directly from the device, reducing latency and shielding back-end name servers from repeated identical queries.
Global Address Database	The device comes with a global address database, providing information such as continent, country, city, operator and organization for IP address queries, used for topology-based GSLB decisions and for geolocation of clients.
Response Policy Zone (RPZ)	Supports import and automatic update of third-party response policy zones to identify domains that are known to be malicious, allowing queries for such domains to be blocked, redirected or walled off at the DNS layer.
Metric and Configuration Exchange	Supports a UDP-based protocol to collect configuration and metric information and exchange that information between load balancing devices. Collection and exchange of configuration and metric information between the supplied devices is secured, with all communication encrypted via SSL.
Object Health Monitoring	Supports monitoring and determining the health and availability of objects — including virtual servers, pools, pool members, links and wide IPs — within the supplied devices, and uses that health state when making GSLB decisions.
Persistence Information Exchange	Supports exchanging virtual server persistence information between the supplied devices running the Local Traffic Manager module, so that a client is returned to the same data centre or virtual server for the duration of a session.
Auto-Discovery	Supports auto-discovery to automatically gather newly created objects, such as virtual servers and links, from the Local Traffic Manager module deployed within the supplied devices, removing the need to define those objects manually in the GSLB configuration.
Synchronization Group Status	Supports exchanging information regarding the status of the DNS synchronization group, so that every member of the group is aware of the availability and configuration state of the other members.
Secure Configuration Synchronization	Provides a secure transport mechanism for synchronizing DNS configuration throughout the synchronization group, so that zone and wide IP changes made on one device are propagated to the others over an authenticated, encrypted channel.

Features	Detail Description
LDNS Path Probing	Supports gathering and communicating local DNS (LDNS) path probing metrics, including round trip time, hop count, packet completion rate and cumulative path score, for use by dynamic GSLB load balancing methods.
GSLB Device Grouping	Supports forming a group between GSLB devices in order to perform automatic configuration synchronization, exchange of server loading information, and exchange of round trip time and hop count information for local DNS servers.
GSLB Load Balancing Algorithms	Supports global load balancing methods based on round robin, global availability, application availability, topology, ratio, bandwidth, round trip time, hop count, completion rate, least connections, packet rate, quality of service and static persistence.
Advanced DNS Statistics	Provides advanced DNS statistical metrics for administrators, including detailed data for profiles such as query type counts with requests, responses and percentage counts.
DNS Reporting and Analytics	Provides advanced DNS reporting and analytics capabilities across applications, virtual servers, query names, query types, client IP addresses and top requested names.
Application Delivery Optimization	
HTTP RAM Cache	Supports HTTP RAM cache. An HTTP cache is a collection of HTTP objects stored in memory that the device reuses to service subsequent requests, configured by associating a Web Acceleration profile with a virtual server.
HTTP/2 Gateway	Supports HTTP/2 and can act as an HTTP/2 gateway, converting HTTP/2 client requests to HTTP 1.1/1.0 on the request sent to the back-end servers, so that existing servers can be front-ended by HTTP/2 without modification.
TCP Connection Multiplexing	Supports TCP connection multiplexing, reusing established server-side connections across multiple client requests to reduce connection setup overhead on the application servers.
Session Persistence in HTTP/2 Full-Proxy	Supports session persistence when operating as an HTTP/2 full proxy. A virtual server with an associated HTTP/2 profile processes the traffic in full-proxy mode while maintaining persistence state.
SSL Offload and Hybrid SSL Acceleration	Offloads SSL/TLS processing to dedicated hardware. The hybrid SSL acceleration feature enables the device to split SSL offload between the hardware accelerator and the CPU for optimal use of available resources.
HTTP Content Modification	Supports rewriting and manipulation of HTTP payload content, allowing retrieval of content, querying of content size and replacement of content in flight.
Maintenance (Sorry) Page	The device can respond with a maintenance page if the virtual server resources (pool or pool members) are unavailable, rather than returning a connection error to the client.
Bandwidth Controller	Depending on the type of policy configured, bandwidth controllers can apply specified rate enforcement to traffic flows or mark traffic for differentiated handling.
X-Forwarded-For Support	When the device translates the source IP address of an incoming packet to a SNAT address, the X-Forwarded-For HTTP header can be inserted to preserve the original client IP address for the benefit of the application server and its logs.
Auto Last Hop	Auto Last Hop allows the device to track the source MAC address of incoming connections and return traffic from pools to the same source, ensuring symmetric return paths in asymmetric or multi-homed topologies.

Features	Detail Description
Application Services Templates	Application Services Templates provide an effective way to deploy applications using AS3 declarative configuration, so that a complete application service can be created from a single declaration.
Management	
SNMP	Supports SNMP v2c and SNMP v3 management protocols, providing a MIB library for device queries and operations, including after device and future version upgrades.
NTP	Supports Network Time Protocol (NTP) for system clock synchronization with one or more external time sources.
Syslog	Supports standard syslog and RSyslog log formats, and allows forwarding of system logs and alarm logs to more than two designated syslog devices, as well as SNMP trap destinations.
Authentication	Supports administrative authentication, including local authentication with locally defined accounts and roles, and remote authentication sources where required.
Dedicated Management Interface	Includes a dedicated out-of-band management interface on each device for GUI-based system management, separate from the data-plane interfaces. Always-On Management (AOM) provides lights-out management through the same Ethernet management port.
GUI and CLI Management	Supports management of the device via both a web-based graphical user interface and a command line interface, accessed over HTTPS and SSH respectively.
Remote Management IP Restriction	Allows restrictions on the IP addresses from which the device may be managed remotely.
System Monitoring and Alerts	Supports real-time monitoring of the device's system resources, including CPU, memory, interface bandwidth, log capacity, number of policies and number of connections, with configurable alerting. Historical CPU utilisation is available from the round robin database (RRD) for trend analysis.
Packet Capture Tool	Provides a real-time packet capture tool that can analyse data packets passing through the device, generating capture files in a standard format for offline analysis.
Statistical Reporting	Provides statistics and analytics charts when Application Visibility and Reporting (AVR) is provisioned, including an overview screen summarising traffic and object health.
sFlow Monitoring	sFlow is an industry-standard technology for monitoring high-speed switched networks. The device can poll internal data sources and send sFlow datagrams to an external collector.
Link Layer Discovery Protocol	Supports Link Layer Discovery Protocol (LLDP), a Layer 2 industry-standard protocol (IEEE 802.1AB) that enables a network device to advertise its identity and capabilities to neighbours.
Platform and Compliance	
Appliance-based Platform	A purpose-built appliance running a single hardened base operating system. The vendor responds to all reported vulnerabilities for that base operating system.

Hardware Specifications

Model: YK-LTM r4800

Specifications	R4800
Intelligent Traffic Processing	L7 requests per second: 1.8M L4 connections per second: 750K L4 HTTP requests per second: 3.5M Maximum L4 concurrent connections: 38M Throughput: 50 Gbps / 40 Gbps L4/L7
Hardware Offload SSL/TLS	45K TPS (2K SSL TPS) 20K TPS (ECDHE-ECDSA P-256 TPS) 20K TPS (ECDHE P-256-RSA 2k TPS) 25 Gbps bulk encryption
Hardware Compression	30 Gbps (Hardware)
Software Architecture	64-bit TMOS 64-bit F5OS
Multi-Tenancy	Up to 4
Processor	1 x 16-Core Intel® Atom® processor (16 CPUs available for Tenancy)
Memory	64 GB DDR
Hard Drive	480 GB M.2 SSD
Management Ports	1 x 1000BASE-T 1 x USB 3.0 1 x serial console
10G/1G Gigabit Copper Ports	4 x 10G/1G RJ45 ports
25/10 Gigabit Fiber Ports	4 x 25G/10G/1G SFP+/SFP28/SFP ports
Power Supply	1 x 250 W 100-240 VAC (+/- 10%) AUTO Switching Platinum, Dual DC PSU optional
Typical Consumption	105 W (single power supply, 48V DC or 110V AC input)
Input Voltage / Typical Heat Generated	Single power supply 110 VAC input: 360 BTU/hour 230 VAC input: 345 BTU/hour
Dimensions	H: 1.74 inches (4.37 cm) x W: 17.1 inches (standard rack) x D: 16.9 inches (42.9 cm). 1U industry standard rack-mount chassis
Weight	20.0 pounds (9.07 kg) with one power supply unit (PSU) (per unit)
Operating Temperature	32 to 104°F (0 to 40°C)
Operational Relative Humidity	5% to 85% (40°C) non-condensing

ETC Information

No.	Feature	Description
(1)	Not exceed 1 RU in height	1U industry standard rack-mount chassis; H 1.74 in (4.37 cm).
(2)	Rack-mountable on standard EIA 19-inch rack	W 17.1 in standard rack; 19-inch rack mount kit supplied.
(3)	At least two hot-swappable AC PSU	Dual hot-swappable AC power supplies supplied as the bid configuration. Base SKU ships with one 250 W PSU; the second PSU is included to meet this requirement.
(4)	At least two internal redundant hot-swappable fan trays, front-to-back airflow	Redundant hot-swappable fan trays with front-to-back airflow.
(5)	At least M.2 480GB SSD local storage	480 GB M.2 SSD.
(6)	At least 64GB DDR memory	64 GB DDR.
(7)	At least four 10/25 (SFP28/SFP+) fiber ports	4 x 25G/10G/1G SFP+/SFP28/SFP ports.
(8)	At least four RJ45 ports	4 x 10G/1G RJ45 ports.
(9)	At least one 100/1000 Mbps Ethernet management port	1 x 1000BASE-T dedicated out-of-band management port.
(10)	At least four 10Gbps multimode SFP+ LC transceivers	Four 10 Gbps multimode SFP+ LC transceivers bundled.
(11)	Active/active and active/passive deployment	Active-Active and Active-Standby redundancy modes.
(12)	Session mirroring for active/active and active/passive	Connection and persistence mirroring in device service clustering, including SSL session mirroring.
(13)	At least 50Gbps Layer-4 throughput	50 Gbps L4 throughput.
(14)	At least 40Gbps Layer-7 throughput	40 Gbps L7 throughput.
(15)	At least 3.0M Layer 4 HTTP requests per second	3.5M L4 HTTP requests per second.
(16)	At least 1.5M Layer 7 requests per second	1.8M L7 requests per second.
(17)	At least 35M Layer 4 concurrent connections	38M maximum L4 concurrent connections.
(18)	At least 25Gbps hardware compression	30 Gbps hardware compression.
(19)	IPv4 and IPv6 GSLB including AAAA resolution	IPv4/IPv6 dual stack GSLB with A and AAAA record resolution.
(20)	Load balancing of DNS servers	DNS server pools with health monitoring.
(21)	Built-in BIND 9 DNS server	BIND 9 based authoritative name server; A, AAAA, CNAME, DNAME, HINFO, MX, NS, TXT, SOA, SRV.
(22)	DNSSEC security protocol	Full DNSSEC support with ZSK/KSK key management.
(23)	DNSSEC zone delegation to secure child zone	Secure delegation with DS record publication.
(24)	Real time DNSSEC signing from authoritative name server	On-the-fly DNSSEC signing of dynamically generated responses.
(25)	Association with EDNS client subnet option	EDNS0 client subnet used for topology decisions.
(26)	Zone transfer from authoritative DNS server	Secondary DNS with AXFR and IXFR zone transfer.

No.	Feature	Description
(27)	DNS caching for authoritative DNS records	Authoritative response cache.
(28)	DNS caching for forwarding DNS queries	Resolver/forwarding cache.
(29)	Source IP address persistence of LDNS	Source IP persistence for LDNS in GSLB.
(30)	Health monitor of devices, applications and content	Active, passive and scripted transaction monitors across devices, applications and content.
(31)	Import and automatic update of third-party RPZ	Response Policy Zone import with automatic feed update.
(32)	UDP-based protocol to collect and exchange config/metrics	UDP-based metric and configuration exchange between devices.
(33)	Secure exchange between H01 devices, SSL encrypted	All inter-device configuration and metric traffic encrypted via SSL.
(34)	Monitoring and determining health of objects within H01	Health determination for virtual servers, pools, members, links and wide IPs.
(35)	Exchanging virtual server persistence information	Persistence state shared between devices running LTM.
(36)	Auto-discovery of objects from LTM module	Automatic discovery of virtual servers and links from LTM.
(37)	Exchange status of DNS synchronization group	Synchronization group status exchange.
(38)	Secure transport for synchronizing DNS configuration	Authenticated, encrypted configuration synchronization.
(39)	Gathering and communicating LDNS path probing metrics	RTT, hop count, completion rate and path score probing.
(40)	Forming group between GSLB devices (sync, load, RTT/hops)	GSLB sync group performing configuration synchronization and exchange of server load and LDNS RTT/hop data.
(41)	Advanced DNS statistical metrics	Per-profile query type counts with requests, responses and percentages.
(42)	Advanced DNS reporting and analytics	Analytics across applications, virtual servers, query names, query types, client IPs and top requested names.
(43)	HTTP RAM cache	Web Acceleration profile providing in-memory HTTP cache.
(44)	HTTP/2 support, acting as gateway to HTTP 1.1/1.0	HTTP/2 gateway converting to HTTP 1.1/1.0 toward back-end servers.
(45)	TCP connection multiplexing	Server-side connection reuse across client requests.
(46)	Session persistence in HTTP/2 full-proxy	Persistence maintained in HTTP/2 full-proxy mode.
(47)	SNMP v2c and v3	SNMP v2c and v3 with MIB library.
(48)	NTP	NTP client with multiple time sources.
(49)	Syslog	Syslog/RSyslog forwarding to more than two destinations.
(50)	Authentication including local authentication	Local accounts and roles; remote authentication also supported.
(51)	Dedicated management interface for GUI management	Dedicated out-of-band 1000BASE-T management port with AOM.
(52)	Management via GUI and CLI	Web GUI over HTTPS and CLI over SSH.