

ZXTM Performance Quick Reference Guide

September 24th 2007

	ZXTM 2000	ZXTM 5000	ZXTM 7000	ZXTM 7400
Layer 4 performance				
Connections / second	11,200	16,200	18,400	33,900
Layer 7 Performance				
HTTP Connections / second	14,800	22,600	26,300	51,100
HTTP Requests / second (0k)	37,900	55,300	64,700	123,000
HTTP Requests / second (2k)	25,000	37,700	43,600	81,300
HTTP Requests / second (8k)	12,200	21,800	24,600	26,600
HTTP max. throughput	0.91 Gbits	2.1 Gbits	2.3 Gbits	2.6 Gbits
SSL transactions per second				
1024 bit RSA, RC4 (no session reuse)	1,940	3,710	4,380	9,420
SSL max. throughput	0.61 Gbits	1.0 Gbits	1.2 Gbits	1.9 Gbits
HTTP Content Caching performance				
HTTP Cache Requests / second	75,600	110,000	129,000	219,000
HTTP Cache max. throughput	1.26 Gbits	3.68 Gbits	3.59 Gbits	3.57 Gbits
HTML and XML Content Transformation performance				
10k HTTP requests/second	n/a	15,600	16,900	21,500
10k HTML content rewrite – rps	n/a	7,910	8,810	13,000
10k HTML content rewrite – rate	n/a	0.64 Gbits	0.77 Gbits	1.1 Gbits
10k XML XSLT – rps	n/a	6,850	8,330	13,700
10k XML XSLT – rate in/out	n/a	0.55 Gbits 0.38 Gbits	0.72 Gbits 0.47 Gbits	1.2 Gbits 0.78 Gbits

These benchmark tests were conducted in Zeus' testing labs using software clients and servers on a dedicated Gbit network. Standard ZXTM 2000, 5000, 7000 and 7400 appliances with the ZXTM 4.2r1 appliance software release were used.

	CPU	Memory	Network Interfaces
ZXTM 2000	Single AMD Opteron 146	2 Gb	2 x 1Gbs, 1 x Mgmt
ZXTM 5000	2 x AMD Opteron 246	4 Gb	4 x 1Gbs, 1 x Mgmt
ZXTM 7000	2 x AMD Opteron 250	4 Gb	4 x 1Gbs, 1 x Mgmt
ZXTM 7400	2 x AMD Opteron 285 (dual core)	8 Gb	4 x 1Gbs, 1 x Mgmt

About the tests

Layer4 Connections per second:

Measures the sustained rate at which new connections can be established from the client to the server, a read-write transaction (HTTP request and zero-byte response body) conducted and the connection closed.

HTTP Connections per second:

Measures the sustained rate at which new connections can be established from the client to the server, an HTTP transaction conducted and the connection closed. This test uses keepalive connections to the server nodes.

HTTP Requests per second:

Measures the sustained rate of HTTP transactions (0Kb, 2Kb and 8Kb response body) using client-side and server-side keepalive connections.

HTTP Cache Requests per second:

Measures the sustained rate of HTTP transactions (0Kb response body) served directly from ZXTM's content cache.

HTML content rewrite:

A 10k html response was captured and rewritten on-the-fly using a TrafficScript rule. Four instances of one string ("media.example.com") were changed to another ("images.example.org"). The two strings were different lengths so this prevented an optimized in-place replacement and properly exercised ZXTM's content rewriting capability.

XML XSLT:

A 10k XML document was returned by the server. This XML document contained 10 search results in the format used by Google's Web Services API. A TrafficScript rule used an XSLT transformation to extract the first search result and return it in HTML format. The result was 339 bytes long, which accounts for the different 'in' and 'out' rate figures.

Performance Considerations

Benchmarks are carefully constructed to extract the maximum performance from the system under test. Real world conditions, such as slow, unreliable network connections, uneven traffic profiles and deep packet inspection mean that benchmark figures (particularly throughput) are unlikely to be achieved when managing live traffic.

If more than 250 Mbits of traffic per ZXTM appliance is expected, or if complex TrafficScript rewriting will be used, performance validation should be conducted. A live deployment should plan for no more than 50% CPU utilization at peak times to accommodate additional traffic spikes and system failover.

IP Transparency incurs additional work and reduces high-volume performance. Peak throughput for large transfers falls by 35%; typical HTTP peak performance falls by 25%; SSL falls by less than 10%.

Real-world Examples

A major content site in the UK operates 4 active ZXTM 7000 systems managing 1.2 Gbits of traffic. Each machine processes around 13,000 requests per second, and runs a number of TrafficScript rules to route requests across 36 backend machines in 3 pools. CPU usage is approximately 6% user, 44% system, 50% idle.

A European-based media service operates 5 active ZXTM software systems (2 processor Dell and HP servers) managing 2 Gbits of traffic at 90% utilisation. Complex TrafficScript rules are in use to manage traffic. They have around 100 web servers.

References

The following whitepapers will be useful:

- Benchmarking Zeus Extensible Traffic Manager.
- Accelerating Web Applications with ZXTM.

Whitepapers are available from <http://www.zeus.com/products/zxtm/>.

The ZXTM KnowledgeHub (<http://knowledgehub.zeus.com/>) makes a number of recommendations on performance tuning a ZXTM system.

