

Advanced Performance Test Report

Cisco

Secure Firewall 4225



ONLINE REPORT

SE LABS tested the performance of the **Cisco Secure Firewall 4225**, assessing its ability to operate under a variety of network loads, including a range of well-established but synthetic sets of traffic and a more realistic mix of protocols.

This test is based on available standards of testing, including the methodology provided by the Internet Engineering Task Force.

The results indicate how effectively the product was at handling network traffic in different circumstances, using the configuration specified.

Contents

Introduction	04
Executive Summary	05
Advanced Performance NGFW Award	05
How we Tested	06
1. Mixed Traffic Capacity Results (variable loads)	07
a. Throughput (Gbps)	07
2. Application Traffic Capacity Results (specific applications)	08
a. Throughput (Gbps)	08
3. HTTP and HTTPS Capacity Results (variable loads)	09
a. Throughput (Gbps)	09
b. Connections/ second	11
c. Transactions/ second	13
4. HTTP and HTTPS Latency Results (normal loads)	15
a. HTTP and HTTPS Connections/ second	15
b. HTTP and HTTPS Transactions/ second	16
5. Conclusions	17
Appendices	18
Appendix A: Device Configuration Details	18
Appendix B: FAQs	19

Document version 1.0 Written 8th May 2025

Management

Chief Executive Officer Simon Edwards

Chief Operations Officer Marc Briggs

Chief Human Resources Officer Magdalena Jurenko

Testing Team

Nikki Albesa

Thomas Bean

Solandra Brewster

Jarred Earlington

Gia Gorbald

Anila Johny

Erica Marotta

Jeremiah Morgan

Julian Owusu-Abrokwa

Joseph Pike

Enejda Torba

Dimitrios Tsarouchas

Marketing

Sara Claridge

Ben Tudor

Publication

Rahat Hussain

Colin Mackleworth

IT Support

Danny King-Smith

Chris Short

Website selabs.uk

Email info@SELabs.uk

LinkedIn www.linkedin.com/company/se-labs/

Blog selabs.uk/blog

Post SE Labs Ltd, 55A High Street, Wimbledon, SW19 5BA, UK

SE Labs is ISO/IEC 27001 : 2022 certified and BS EN ISO 9001 : 2015 certified for The Provision of IT Security Product Testing.

SE Labs is a member of the Microsoft Virus Initiative (MVI); the Anti-Malware Testing Standards Organization (AMTSO); the Association of anti Virus Asia Researchers (AVAR); and NetSecOPEN.

© 2025 SE Labs Ltd

Introduction



CEO
Simon Edwards

If you spot a detail in this report that you don't understand, or would like to discuss, please contact us. SE Labs uses current threat intelligence to make our tests as realistic as possible. To learn more about how we test, how we define 'threat intelligence' and how we use it to improve our tests please visit our [website](#) and follow us on [LinkedIn](#).

Standards, Analysis and Transparency

This network performance test is designed to help you make the most informed buying decisions.

This test was conducted using recommendations made by the Internet Engineering Task Force for testing the performance of network security devices such as next-generation firewalls, intrusion detection and protection systems and unified threat management devices. At a minimum it should show, in a transparent and repeatable way, how the device under test handles network traffic of different types and in different scenarios.

On its own, the raw data is useful for comparing products with a view to choosing which is most suitable for your organisation.

We also ran an extended set of tests to see how a device would behave in a more realistic, production environment. This involved using a mixture of network traffic protocols and testing individual types of application traffic.

At SE Labs we don't just publish raw figures, though. We use our knowledge and expertise to analyse that information to help add useful colour to the results.

The goal is to give a real-world opinion as to which figures are most important, highlight where optimum performances are achieved and to explain why some details are more significant than others.

For example, a device might achieve an apparently strong performance when handling Voice over IP, but in real-life the human ear might struggle with sub-par connection quality. Conversely, what may seem like poor performance on paper might not be noticeable to users in a real deployment.

We have followed the available testing standards so that you can verify our figures with reports generated by other test labs. This gives you confidence that the testing was conducted correctly, while also being completely transparent about the configuration used. This configuration might not be the one you experience out of the box or might not be suitable for your own deployment. We've included configuration details so you can make a fully informed decision when comparing products and reports.

Executive Summary

This test assesses the product's ability to handle different levels of network traffic while its security features are enabled.

It includes assessments made under a variety of network loads, including a range of well-established but synthetic sets of traffic and a more realistic mix of protocols. This test is based on available standards of testing, including the methodology provided by the Internet Engineering Task Force.

The results indicate how effectively the product was at handling network traffic in different circumstances, using the configuration specified.

- The 30Gbps-rated Cisco Secure Firewall 4225 was able to handle realistic traffic loads optimally.
- Different applications were handled at rates ranging from around 47Gbps and 8Gbps. FTP and SMB protocols were handled most effectively. Overall this falls below the benchmark target of 24Gbps.
- The device handled web-based traffic at rate of 57Gbps for HTTP, which is above the minimum requirement of 30Gbps. The maximum throughput rate for encrypted traffic was over 30Gbps, which exceeds the target of 15Gbps.


Advanced Performance NGFW Award

The following product wins the SE Labs award:

Performance Summary

Test	Status
Mixed Traffic Capacity	PASS
HTTP Capacity	PASS
HTTPS Capacity	PASS
HTTP/S Latency (CPS) Overall	PASS
HTTP/S Latency (TPS) Overall	PASS

- Web traffic latency results, which indicate how responsive users would find using the web, were excellent for both unencrypted and encrypted traffic.



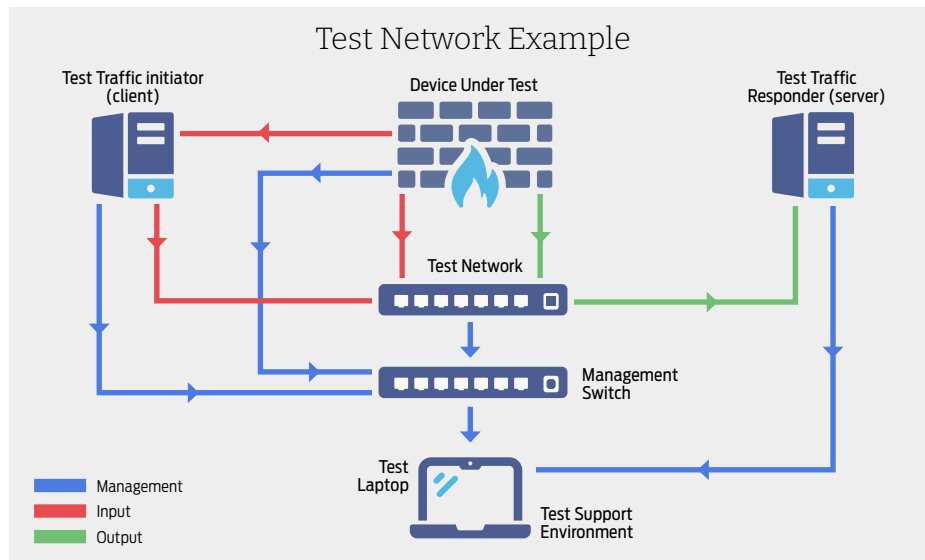
How We Tested

There is a lot to consider when choosing a network security appliance, and its speed performance on the network is a significant factor. This set of tests is designed to give a good idea how well the device being tested can perform in realistic, production environments as well as exploring the edgier laboratory conditions possibilities.

The results cover how quickly the device can shift different types of network traffic. We tested using a realistic mixture of enterprise traffic, specific applications and network services and moved on to examining detailed results for throughput and latency.

Throughput tests show how much data can pass through the device before it becomes overwhelmed and slows things down. But that's just one part of the story. Latency, which indicates how responsive users will find their experience on the network, is also critical to a productive deployment. For this reason we measured latency, and in more than one way. We looked at how fast web pages can be downloaded in full, and how quickly users can expect to see a connection at least start.

When testing the network performance of a network device there needs to be a network traffic load to either demonstrate or push it beyond its abilities. The more realistic this load, the more realistic and (therefore) useful the test is.



We used the load details specified by the Benchmarking Methodology Working Group of the Internet Engineering Task Force, which is supported by the NetSecOPEN standards organisation. Specifically, we followed version 06 of the draft document [Benchmarking Methodology for Network Security Device Performance V6](#).

We used a variety of load specifications using the Spirent CyberFlood Virtual. We used Spirent's SimUsers/Second feature to generate traffic in all but the Connections per Second test.

Test set-up

The configuration of the device was based on recommendations by its vendor. Details are available in **Appendix A: Device Configuration Details** on page 18.

1. Mixed Traffic Capacity Results

This test indicates how much data can pass through the device in a real-world production environment, rather than a sterile and theoretical laboratory test. It should answer the question, “how much throughput can I expect if I buy and use this?”

A realistic mixture of network traffic using different protocols, as might be expected to pass through an enterprise network firewall, is sent through the device.

This challenging test requires the device to check, track and respond to lots of different types of connections. It shows devices at their best or worst.

We consider the optimum throughput result to be at least 50% of the device’s stated maximum. Anything above 75% is excellent.


We highlight below the point at which the device’s performance most closely matches the traffic load.

Traffic Mix

Protocol/Application	Percentages
HTTPS 1.1	25.08%
HTTP 1.1	15.08%
SMTP	12.12%
IMAP4	12.08%
SMBv2	10.08%
Oracle	8.08%
FTP	7.58%
MySQL	4.08%
RTSP	4.08%
Syslog	1.18%
RDP	0.28%
SSH	0.28%

Mixed Traffic Tests

Throughput (Gbps)	16Gbps	17Gbps	18Gbps	19Gbps	20Gbps	21Gbps	22Gbps	23Gbps	24Gbps	25Gbps	26Gbps	27Gbps
Throughput Achieved (Gbps)	16	17	18	18.9	19.9	20.9	21.9	22.9	23.9	25	26	27
Average URL Response Time (ms)	57.4	60.7	65.7	68.4	73.6	75.9	82.1	95.3	110.9	116.6	2,350.7	2,736
Unsuccessful Transactions	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	2.3%	2.8%



2. Application Traffic Capacity Results

This test indicates how well the device can perform with specific applications and services, rather than the mixture of protocols used in the Mixed Traffic Capacity test. The applications and services were tested in isolation to each other and not concurrently.

The results should highlight specific strengths and weaknesses in the device's ability to handle different types of network traffic. For example, a

device might achieve a high throughput for FTP traffic, but SMTP traffic performance could be lower. SMB throughput might be high, but Skype might suffer due to a lower throughput, introducing latency issues that are important to avoid when video conferencing.


We tested with different loads until errors reached a threshold of over 1%. We then reported on the previous load, achieved before the error rate reached that threshold.

We consider the optimum throughput result to be 80% or higher of the maximum load.

The device should achieve this with all applications and to pass.

Application Specific Throughput

Application	SIP	SMTP	FTP	SMB	FIX	RDP	Exchange	Skype	Oracle
Throughput (Gbps)	9.1	16.9	46.6	32.9	7.71	12.8	7.8	27.3	23.2




3. HTTP and HTTPS Capacity Results

These tests indicate how much data can pass through the device when it is handling web sessions.

It submits the tested device to a range of loads of mixed body sizes, starting with a low amount of network traffic and measuring its ability to transfer data as that load increases. The throughput


measurements show how busy a network can be before the device starts to struggle and under-perform. At the same time, the test measures how many connections and transactions per second are possible.

We consider the optimum throughput result for HTTP to be when the device successfully processed the highest

load without slowing the data transfer. For example, if it can receive 5Gbps of traffic and transfer this data through at 5Gbps, then that's an acceptable result. However, if it transfers 6Gbps of data in just 5.5Gbps then it's slowing the overall transfer of data. We expect HTTPS traffic to run at around half-speed due to the overhead encryption imposes. 

HTTP Throughput

HTTP Load	35Gbps	40Gbps	45Gbps	50Gbps	55Gbps	56Gbps	57Gbps	58Gbps	60Gbps	65Gbps	70Gbps	75Gbps	80Gbps
Throughput Achieved (Gbps)	36	41.1	46.3	51	56.6	57.6	56	36.6	3.2	1.8	1.3	1.4	1.2
Average URL Response Time (ms)	2.15	3.28	3	5.2	9.8	6.3	14.3	111.5	427	436.6	644.9	1,131.6	580.3
Unsuccessful Transactions (%)	0%	0%	0%	0%	0%	0%	1.7%	1.9%	15.4%	13.4%	15.2%	16.7%	18.4%



We measure throughput speeds in increments of 5Gbps until the device starts to reach its limits. We then switch to 1Gbps increments to show a more detailed set of data as the device works up to and past its useful limit.


To pass this test the device should be able to transfer HTTP traffic without slowing it down, at speeds of 80% or more of its stated maximum speed.

For example, if the device is rated at 10Gbps then it should be able to allow an 8Gbps HTTP traffic load to pass through at 8Gbps. A 10Gbps HTTPs load should pass through no slower than 4Gbps.

We highlight below the point at which the device's performance most closely matches the traffic load at the highest load.

HTTPS Throughput

HTTPS Load	5Gbps	10Gbps	15Gbps	20Gbps	25Gbps	30Gbps	31Gbps	32Gbps
Throughput Achieved (Gbps)	5.2	10.3	15.5	20.6	25.8	30.9	32	33
Average URL Response Time (ms)	1.2	1.5	1.9	2.4	3	4.3	5.5	5.6
Unsuccessful Transactions (%)	0%	0%	0%	0%	0%	0%	0%	0%



Each bar's height shows the throughput achieved by the device when loaded with increasing levels of traffic.

We measure throughput speeds in increments of 5Gbps until the device starts to reach its limits. We then switch to 1Gbps increments to show a more detailed set of data as the device works up to and past its useful limit.

HTTP and HTTPS Connections/Second (CPS)

This test loads the device with increasing numbers of web connections to see how it handles different ranges of use.

The Connections Per Second (CPS) measurements show how many basic web connections are possible at any one time. A basic request to a web server, and its response, is a connection.

For example, requesting and receiving a single HTML page counts as one connection.


We test by sending thousands of such requests and measure how many responses the device allows. We also measure the how long it takes for the conversation (the request and the response) to complete. This is the Average URL

Response Time, which is measured in milliseconds. A fast response means a snappy user experience when browsing the web. We define 'fast' as being under 1.5ms.

We expect the CPS achieved to closely match the CPS load when testing using HTTP. HTTPS responses may be half that.


HTTP Connections/Second (CPS)

HTTP CPS Load	120,000 CPS	122,000 CPS	124,000 CPS	126,000 CPS	128,000 CPS	130,000 CPS	132,000 CPS	134,000 CPS	136,000 CPS	138,000 CPS
CPS Achieved	120,522	122,674	124,878	126,944	129,221	131,132	133,236	135,882	135,661	139,745
Average URL Response Time (ms)	1.7	1.8	1.9	2	2.3	2.2	2.4	3.6	3.3	3.6
Unsuccessful Transactions (%)	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%



HTTPS Connections/Second (CPS)

HTTPS CPS Load	2,000 CPS	4,000 CPS	6,000 CPS	8,000 CPS	10,000 CPS	12,000 CPS	14,000 CPS	16,000 CPS	18,000 CPS	20,000 CPS
CPS Achieved	2,048	4,125	6,043	8,108	10,141	12,093	14,088	16,043	18,039.00	20,024.00
Average URL Response Time (ms)	1.4	1.4	1.4	1.6	1.8	2	2.3	2.5	2.9	3.7
Unsuccessful Transactions (%)	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%



Each bar's height shows the throughput achieved by the device when loaded with increasing levels of traffic.

HTTP and HTTPS Transactions/Second (TPS)


The Transactions Per Second (TPS) measurements show how many groups of basic web connections are possible at any one time. A group of connections means multiple requests and responses such as you would experience when loading a web page containing text, images and other elements. For example, an HTML page, some images and an audio file.

We test by sending thousands of such groups of requests and measure how many responses the device allows. We also measure how long it takes for the conversation (the requests and responses) to complete. This is the Average URL Response Time, which is measured in milliseconds. A fast response means a snappy user experience when browsing the web. We define 'fast' as being under 1.5ms.

We expect the TPS achieved to closely match the TPS load when testing using HTTP. HTTPS responses may be half that.


HTTP Transactions/Second (TPS)

HTTP TPS Load	170,000 TPS	180,000 TPS	190,000 TPS	200,000 TPS	210,000 TPS	220,000 TPS	230,000 TPS	240,000 TPS	250,000 TPS	260,000 TPS
TPS Achieved	169,869	179,919	189,824	199,787	209,779	219,768	229,747	239,671	44,617	32,087
Average URL Response Time (ms)	3.7	3.6	6.3	5.6	6.3	5.8	9.8	11.7	663.3	1,073.7
Unsuccessful Transactions (%)	0%	0%	0%	0%	0%	0%	0%	0%	14.9%	16.2%



HTTPS Transactions/Second (TPS)

HTTPS TPS Load	40,000 TPS	50,000 TPS	60,000 TPS	70,000 TPS	80,000 TPS	90,000 TPS	100,000 TPS	110,000 TPS	120,000 TPS	130,000 TPS
TPS Achieved	40,151	50,163	60,161	70,161	80,513	90,114	100,083	110,028	119,985	129,977
Average URL Response Time (ms)	1.6	1.7	1.9	2.1	2.7	2.9	3.2	4	4.3	5.9
Unsuccessful Transactions (%)	0%	0%	0%	0%	0%	0%	0%	0%	14.9%	16.2%



4. HTTP and HTTPS Latency Results

This test indicates how responsive the device is when operating under normal loads.

Latency is measured in two ways: by timing how long it takes to download the full body of a transaction (e.g. a web page) and by timing how long it takes for the first piece of the web page to be received by the client's browser. The results are called 'URL response time' and 'time to first data byte' respectively.

Together the latency measurements show how smoothly users experience web browsing when the device is on the network. The URL response time shows how quickly they can expect to download full pages, while the 'time to first data byte' results shows how fast they experience the beginning of a connection.


For example, a fast 'time to first data byte' result would mean that the user would see the web

browser connect fast to the website and start downloading content. However, a relatively slow URL response time would mean that the page itself, and the elements it contains, might take some time to download fully. In contrast, a slow 'time to first data byte' result would mean that the user waits for the initial connection to establish but, if the URL response time was fast, the page would then quickly download.




HTTP Latency Connections per Second

	Transaction Body Sizes		
HTTP Connections per Second	1KB	16KB	64KB
CPS Achieved	120,920	65,865	35,792
URL Latency (average URL response time (ms))	0.61	0.65	1.02
Average Time To First Data Byte (ms)	0.61	0.55	0.66



HTTPS Latency Connections per Second

	Transaction Body Sizes		
HTTPS Connections per Second	1KB	16KB	64KB
CPS Achieved	11,205	10,681	10,109
URL Latency (average URL response time (ms))	0.88	1.53	3.08
Average Time To First Data Byte (ms)	0.88	1.53	1.56



We take an average value from the HTTP and HTTPS results to create an overall result. To pass this test the device should achieve an average latency result of below 2ms.

Good results are under 2.0ms to first data byte.
An excellent result is under 1.5ms to first data byte.

In our tables in this section, Connections Per Second show how many times the device could handle the creation of a connection, a single transaction (a webpage download) and then the termination of the connection. Transactions Per

Second is the same, but with 10 transactions made between the start and end of the connection.

Test Load Details

The way we set up the load for this test is complicated and based on a number of factors, including the requirements of the IETF's Benchmarking Methodology for Network Security Device Performance (RFC 9411).


Essentially, testers must establish the maximum number of connections and transactions the device can handle every second, without significant error

levels, and then test with loads of half those sizes. Unlike the throughput tests, the loads are not measured in Gbps but in CPS and TPS.

For example, if the tester finds that the device can reliably handle 20,000 Connections Per Second (CPS), the test should use a load of 10,000 CPS. The same goes for the Transactions Per Second (TPS) test: find the reliable maximum and then test with a load of half that size. We run the tests using three different data sizes: 1Kb, 16Kb and 64Kb. These are the figures you see here and above.


HTTP Latency Transactions per Second

	Transaction Body Sizes		
HTTP Transactions per Second	1KB	16KB	64KB
TPS Achieved	249,272	149,736	60,164
URL Latency (average URL response time (ms))	0.56	0.79	1.68
Average Time To First Data Byte (ms)	0.56	0.55	0.61



HTTPS Latency Transactions per Second

	Transaction Body Sizes		
HTTPS Transactions per Second	1KB	16KB	64KB
TPS Achieved	99,966	78,100	40,207
URL Latency (average URL response time (ms))	0.9	1.5	3.36
Average Time To First Data Byte (ms)	0.92	1.5	1.59



5. Conclusion

This test assesses the product's abilities to handle different levels of network traffic while its security features are enabled.

Network security appliances are designed to achieve two main goals: to allow legitimate traffic to pass through the network unhindered and to apply security controls that handle unwanted traffic. They may also prioritise certain types of traffic over others, improving performance where it will be most noticed by the organisation using it.

To examine the product's performance, we looked at four main areas. Potentially the most interesting is where the device is loaded with realistic network traffic to see how closely its throughput matches that claimed by its manufacturer.

Manufacturers claim different performance levels depending on how the firewall is configured. We used the security settings listed in the Features section of **Appendix A: System Configuration Details**.

Cisco claims a throughput rate of 30Gbps for the **Cisco Secure Firewall 4225** when configured this way.

Realistic Network Loads

This part of the test should answer the question, "how much throughput can I expect if I buy and use this?" We consider the optimum throughput

result to be at least 50% of the device's stated maximum. Anything above 75% is excellent.

In this case, **Cisco Secure Firewall 4225's** maximum is supposed to be 30Gbps, so a good result would be anything equal to or exceeding 15Gbps. The results seem to indicate that the firewall was capable of matching network loads up to 25Gbps. And while it was able to move more, the error rates increased, as did latency.

The URL response times were good up to this point, which means that users browsing the web through this device can expect a smooth experience. An excellent result.

Result: **PASS**

Application-Specific Loads

The Application Traffic Capacity Results show how the product handles specific individual applications and services, such as the SMTP email protocol and Skype service traffic.

We consider the optimum throughput to be 80% or more of the maximum load. Since **Cisco** claims a maximum throughput rate of 30Gbps for the **Cisco Secure Firewall 4225**, a good result would be anything equal to or exceeding 24Gbps.

The fastest handling of the applications tested was for FTP and SMB, which the device allowed through

at a speed of 46.6Gbps and 32.9Gbps respectively. These results are above the benchmark of 24Gbps.

Handling of Skype traffic topped out at 27.3Gbps which is a good result. SMTP traffic ran through at 16.9Gbps, which is lower than our expectations. There were much lower results for RDP, Exchange, SIP, FIX and Oracle.

Web Traffic Throughput

The HTTP and HTTPS Capacity Results indicate how effective a device is at handling web-based traffic, both encrypted and unencrypted. We submit the device to a range of loads, starting with a low amount of traffic, and measure its ability to transfer data as that load increases. We also measure the connections and transactions per second. Ideally a device will process the load without slowing it, up to its stated capacity.

The Cisco Secure Firewall 4225 is rated at 30Gbps and so should be capable of handling 30Gbps of HTTP unencrypted web traffic and up to 15Gbps of encrypted data without slowing the flow.

The firewall allowed 57Gbps of unencrypted traffic, which exceeds the stated target of 30Gbps for HTTP loads. Above this load the device's performance dropped to extremely low levels of around 1Gbps.

The firewall should process encrypted web traffic no slower than 15Gbps. It achieved this goal, managing to handle up to 30Gbps albeit with increasing amounts of latency. At loads of 31-32Gbps it increased latency but still managed to push the data through at the correct speed.

Result: **PASS**

Web Traffic Latency

The Latency Results indicate how responsive a product is when put under a 'normal' load, that being defined by the IETF as 50% of the maximum throughput achieved in the HTTP and HTTPS Capacity Results (see above). We consider an optimum result here to be 2ms or under for the 'time to first data byte' measurements.

The **Cisco Secure Firewall 4225** achieved results of between 0.55ms and 0.66ms for Connections Per Second with HTTP loads. These are excellent results that do not solely depend on transaction size. It took an average of 0.61ms to get the first data byte of a 1Kb load compared to an average of 0.66ms to that of a 64Kb load.

The HTTPS results for Connections Per Second were also very good, ranging from 0.88ms to 1.56ms.

The Transactions Per Second latency results for HTTP were between 0.55ms to 0.61ms, which is also excellent. The HTTPS transactions were inevitably slower, with latencies between 0.92ms to 1.59ms.

Result: **PASS**

Last Words

The 30Gbps-rated **Cisco Secure Firewall 4225** was able to handle realistic traffic loads optimally, managing 25Gbps. In pure HTTP and HTTPS throughput tests things looked even stronger, moving HTTP traffic at 57Gbps and HTTPS traffic at 32Gbps. Web traffic latency was uniformly excellent for both HTTP and HTTPS traffic.

Enterprise Security Testing Services for CISOs

Elevate your cyber security strategy with SE Labs, the world's leading security testing organisation.

SE Labs works with large organisations to support CISOs and their security teams:

- Validate existing combination of security products and services.
- Provide expert partnership when choosing and deploying new security technologies.

SE Labs provides in-depth evaluations of the cyber security that you are considering, tailored to the exact, unique requirements of your business.

For an honest, objective and well-informed view of the cyber security industry contact us now at

selabs.uk/contact

Appendices

Appendix A: System Configuration Details

Device Under Test Details

Make and Model	Cisco Secure Firewall 4225
NGFW Version	7.6.0
Serial Number	FJZ27491E1N
Snort Version	2.9.23 (build 227)
Snort3 Version	3.1.79.1 (Build 121)
Rule Pack Version	3060
Module Pack Version	3444
LSP Version	lsp-rel-2024111113-1921
VDB Version	Build 397 (2024-10-08 17:51:44)
Rule Update Version	2024-11-13-001-vrt
Geolocation Version	Country Code:2024-11-09-057, IP:None
OS	Cisco Firepower Extensible Operating System (FX-OS) 2.16.0 (build 128)
Vendor Throughput Rating (security settings)	80Gbps

Network Details

	Management Interface	Client Interface	Server Interface
Interface	Mgmt	Eth3-1	Eth3-2
Physical Interface	Copper 10/100/1,000	Fibre 100,000	Fibre 100,000
Physical configuration	Auto	Auto	Auto
Zone	Management	Client	Server

Features

Features Enabled	Enabled/Disabled
SSL Inspection	Enabled
IDS/IPS	Enabled
Threat Defense Malware Protection	Partially Enabled
Advanced Malware Protection	Enabled
Logging and Reporting	Enabled
Application Visibility and Control	Enabled
Security Intelligence	Disabled

Test Equipment Details

Make and Model	Spirent CFv
Controller Software Version	25.1.1008
Firmware Version	25.1.7011

Appendix B: FAQs

Q What is a partner organisation? Can I become one to gain access to the threat data used in your tests?

A Partner organisations benefit from our consultancy services after a test has been run. Partners may gain access to low-level data that can be useful in product improvement initiatives and have permission to use award logos, where appropriate, for marketing purposes. We do not share data on one partner with other partners. We do not partner with organisations that do not engage in our testing.

Q I am a security vendor and you tested my product without permission. May I access the threat data to verify that your results are accurate?

A We are willing to share a certain level of test data with non-partner participants for free. The intention is to provide sufficient data to demonstrate that the results are accurate. For more in-depth data suitable for product improvement purposes we recommend becoming a partner.

Q How do you score awards?

A We add up how many tests the product passes and allocate an award based on the table below:

Awards

Award	Criteria
AAA	Excellent in all five test elements
AA	Excellent in at least three of the test elements; Good in the remainder
A	Good in all five test elements
B	Good or Excellent in three test elements
C	Good or Excellent in two test elements

To pass each test the device must achieve certain criteria, as listed below. These are measured against the device's stated maximum speed, which is a value claimed by the manufacturer. This value depends on the device's configuration. We list both the vendor's stated maximum throughput rating and the security settings that were enabled in **Appendix A: System Configuration Details**.

Test	Pass Criteria (Excellent)	Pass Criteria (Good)
1. Mixed Traffic Capacity	75-100%	50-75%
3. HTTP Capacity Results	90-100%	80-90%
4. HTTPS Capacity Results*	90-100%	80-90%
5. HTTP/S Latency (CPS)*	0-1.5ms	1.5-2ms
6. HTTP/S Latency (TPS)*	0-1.5ms	1.5-2ms

* HTTPS is measured against 50% of the device's stated maximum speed.

A **full methodology** for this test is available from our website.

- The product chosen for this test was selected by SE Labs.
- The test was sponsored Cisco, Inc.
- The test was conducted between 18th March and 3rd April 2025.
- The product was configured according to its vendor's recommendations.
- Test data was provided to partner organisations once the test was complete.

SE Labs Report Disclaimer

1. The information contained in this report is subject to change and revision by SE Labs without notice.
2. SE Labs is under no obligation to update this report at any time.
3. SE Labs believes that the information contained within this report is accurate and reliable at the time of its publication, which can be found at the bottom of the contents page, but SE Labs does not guarantee this in any way.
4. All use of and any reliance on this report, or any information contained within this report, is solely at your own risk. SE Labs shall not be liable or responsible for any loss of profit (whether incurred directly or indirectly), any loss of goodwill or business reputation, any loss of data suffered, pure economic loss, cost of procurement of substitute goods or services, or other intangible loss, or any indirect, incidental, special or consequential loss, costs, damages, charges or expenses or exemplary damages arising his report in any way whatsoever.
5. The contents of this report does not constitute a recommendation, guarantee, endorsement or otherwise of any of the products listed, mentioned or tested.
6. The testing and subsequent results do not guarantee that there are no errors in the products, or that you will achieve the same or similar results. SE Labs does not guarantee in any way that the products will meet your expectations, requirements, specifications or needs.
7. Any trade marks, trade names, logos or images used in this report are the trade marks, trade names, logos or images of their respective owners.
8. The contents of this report are provided on an "AS IS" basis and accordingly SE Labs does not make any express or implied warranty or representation concerning its accuracy or completeness.