



Avalanche Commander FAQs

This document provides the following Frequently Asked Questions (FAQs) for Avalanche Commander:

Setup and Configuration FAQs	3
Are there different TGZ files?.....	3
How do Analyzer, Excel, and Acrobat Reader work with Avalanche?.....	3
How do I back up my tests?	3
How do I check the serial number on a SmartBits unit?	3
How do I configure my Avalanche test ports to auto-negotiate?	3
How do I convert a test from Network Edition 5.2 or Enterprise Edition 6.1 or earlier?.....	3
How do I convert a test from Network Edition 6.2?	3
How do I export and import tests?.....	4
How do I find and modify my network settings?.....	4
How do I find the total time my test will run?.....	4
How do I find what type of license I have?.....	4
How do I get my test to produce a PCAP file, and how do I view PCAP files?	4
How do I perform an OS upgrade?.....	4
How many user profiles and subnets are supported?	4
How do I revert back to factory default settings?	4
How do I use 1s and 2s in a URL list?.....	5
Should I use HTTP 1.0 or HTTP 1.1?	5
What is Keep Alive/Persistence?	5
What is Passive FTP?	5
What is a virtual router and how do I use one?.....	5
Is it possible to automate tests by using the WorkSuite Manager or Tcl API for Windows?..	6
Troubleshooting Hardware and Software FAQs	6
How do I find error messages?.....	6
What are common error messages found in the Event Log files and what do they mean?..	6
How do I find system logs?.....	6
What TCP stack does Avalanche use?.....	6
What is the DNS timeout on Avalanche?.....	7
What is the OS base?	7
What is the Port Mapping for Avalanche?	7

What is the difference between a simuser, connection, and transaction?	7
What should I do if my test produces failures?	8
What types of cables are needed for my hardware?.....	8
When I run my test, why does the interface keep ARPing instead of sending traffic?	9
Why can't my test connect to the SmartBits chassis when I run my test?	9
Why do I see resets at the end of my tests?	9
Why does Avalanche not respond in Commander, even though I'm able to Telnet into my hardware?	9
Why should I use sample tests?	9
How do I troubleshoot FTP problems such as throttling down, instability, and rejected connections?	9
How do I troubleshoot unsuccessful streaming session?.....	10
Why don't I get better performance when I use more than one port with a test (assuming the ports are on the same card)?	10
Why do my tests terminate prematurely when using SmartBits?	10
Why do my performance tests behave differently from version 5.2 to 7.0?	10
Why can other users access my appliance while my test is running, causing it to fail?	10
Why would my appliance become inaccessible through FTP and Telnet?	10
Why do I get RX Underrun/Overrun errors in my results summary CSV file?	11
What do I use for my installation password?	11
Contacting Spirent Communications	12

Setup and Configuration FAQs

Are there different TGZ files?

There is only one TGZ file for both Avalanche and Reflector.

How do Analyzer, Excel, and Acrobat Reader work with Avalanche?

Use Analyzer and Microsoft Excel to view your test results. Analyzer is installed with Avalanche Commander and presents your test result files graphically.

Use Microsoft Excel (2000 or higher) to view result files in a spreadsheet format. You can also use Excel to create form database files or Action lists.

Use Acrobat Reader (4.0 or higher) to view online guides.

How do I back up my tests?

You can export a project and one or more of its tests to create a backup.

1. Select **File > Export Projects**.
2. Follow the instructions that appear in the *Export Tests* wizard.

How do I check the serial number on a SmartBits unit?

1. Select **Administration > SmartBits Chassis**.
2. Select **Chassis > View Chassis Info**.

How do I configure my Avalanche test ports to auto-negotiate?

Your hardware is set to automatically auto-negotiate. Appliances do not have an option to force duplex and work best against devices that are also set to auto-negotiate.

For SmartBits, set the ports to force or auto-negotiate:

1. Select **Administration > SmartBits Chassis**.
2. Select the desired setting from the Auto Negotiation section of the Ports pane.

How do I convert a test from Network Edition 5.2 or Enterprise Edition 6.1 or earlier?

1. It's recommended that you first create a backup. Contact [Spirent Communications Technical Support](#) for assistance when backing up files before you upgrade.
2. Select **File > Import Legacy Tests**.
3. Follow the instructions that appear in the *Import Test* wizard.

NOTE: Avalanche and Reflector tests are combined into a single test file.

How do I convert a test from Network Edition 6.2?

1. It's recommended that you first create a backup. Contact [Spirent Communications Technical Support](#) for assistance when backing up files before you upgrade.
2. Using Avalanche Commander 6.2, select **File > Export Projects**.
3. Using your latest version of Commander, select **File > Import from Archive**.
4. Locate the file you created in Step 2, and then click **OK**.

How do I export and import tests?

You cannot import or export individual tests. Export and import the project in which the test resides.

To export, see [“How do I back up my tests?”](#)

To import a project and tests that were created by using Avalanche 6.2 or later, see [“How do I convert a test from Network Edition 6.2?”](#)

To import a project and tests that were created by using Network Edition 5.1 or Enterprise, see [“How do I convert a test from Network Edition 5.2 or Enterprise Edition 6.1 or earlier?”](#)

How do I find and modify my network settings?

For SmartBits: Select **Administration > SmartBits Chassis**.

For Appliances: Select **Administration > Appliances**, and then click the **Network Configuration** tab.

How do I find the total time my test will run?

Click the **Client** tab, and then the **Loads** tab. The Timing pane shows the total duration of the test based on your load settings.

How do I find what type of license I have?

Select **Administration > View License**.

How do I get my test to produce a PCAP file, and how do I view PCAP files?

Perform a Trial Run to produce a PCAP file:

1. Select the test from the navigation tree.
2. Click the **Trial Run** icon.
3. To view the file, click the **Results** tab.

To produce a PCAP file in addition to result files from a full test run:

1. Select the test from the navigation tree.
2. Select the **Enable Server Packet Trace** or **Enable Client Packet Trace** checkbox from the Run Configure tab. The PCAP file shows all packets sent and received, and can be read by many standard Sniffer programs, such as Ethereal, Netasyst, Sniffer Pro, or ClearSight.

How do I perform an OS upgrade?

Contact [Spirent Communications Technical Support](#).

How many user profiles and subnets are supported?

Avalanche Commander supports up to 100 user profiles per subnet, and up to 250 subnets per test port. The number of test ports depends on your hardware.

How do I revert back to factory default settings?

Contact [Spirent Communications Technical Support](#).

How do I use 1s and 2s in a URL list?

The URL list indicates the level and method of access or command used.

- URL entries that start with a 1 emulate a URL explicitly clicked by a user such as an index.html.
- URL entries that start with a 2 emulate a URL embedded in a level-1 URL retrieved by a browser, such as a .jpg or .gif object.

Entries that specify only the level-1 URL, retrieve only the main index.html file text without any embedded objects, such as banners and graphics.

For example, the following retrieves only the main index.html file text without any embedded objects, such as banners and graphics.

```
1 http://www.somewebsite.com/index.html
```

To retrieve the entire page, specify all the URL elements, such as

```
1 http://www.somewebsite.com/index.html
```

```
2 http://www.somewebsite.com/logo.gif
```

```
2 http://www.somewebsite.com/banner.jpg
```

Level 2s are only valid in HTTP/HTTPS tests and are ignored for other protocols. For information, see "HTTP Action Lists" in the Avalanche Commander Help.

Should I use HTTP 1.0 or HTTP 1.1?

This depends on the goal of the test. HTTP 1.0 without Keep Alive results in one TCP connection per HTTP transaction, whereas HTTP 1.1 with Persistence results in many HTTP transactions (such as GETs) within a single TCP connection. In general, if the goal of the test is only concerned with connections, then HTTP 1.0 without Keep Alive is suggested. If the testing goal is HTTP transactions, then HTTP 1.1 with Persistence is suggested.

See the "HTTP Test Scenarios" in the Avalanche Commander Help that illustrate the client and server behavior as a result of various client Protocol Level Field parameter settings, and server Connection Properties and Server Emulation parameter settings.

What is Keep Alive/Persistence?

Keep Alive, also known as Persistence, allows multi HTTP 1.1 transactions (Posts or GETs) to be performed within a single TCP connection by keeping the connection open. This is the default behavior of browsers such as Internet Explorer (Netscape uses HTTP 1.0 with Keep Alive).

What is Passive FTP?

In passive mode, the FTP client initiates both the control connection and data link connection to the server. When opening an FTP connection, the client opens two random unprivileged ports locally ($N > 1024$ and $N+1$). The first port contacts the server on port 21, but instead of then issuing a PORT command and allowing the server to connect back to its data port, the client issues the PASV command. The server then opens a random unprivileged port ($P > 1024$) and sends the PORT P command back to the client. The client initiates the connection from port $N+1$ to port P on the server to transfer data. Only passive FTP is supported on Avalanche Commander.

What is a virtual router and how do I use one?

A virtual router emulates a router at the test port interface, and sets up a primary IP address through which all requests going into and coming out of the test port will be passed through a single IP and MAC address. This allows you to emulate one or more unique client subnets behind the virtual router. The virtual router IP address needs to be on the same subnet as the gateway or router to which it is connected. If you don't use a virtual router,

then your client subnet must be on the same network as your directly connected device under test. Either the virtual router or a static route should be selected as needed.

Is it possible to automate tests by using the WorkSuite Manager or Tcl API for Windows?

You can use the WorkSuite Manager to create WorkSuites, a collection of Avalanche Commander tests that you can execute consecutively and unattended. For more information, see the Avalanche Commander Help.

You can automate tests by using the Tcl API instead of using the Avalanche GUI. Basically, the Tcl API enables you to write scripts that perform the same actions as the Commander GUI. For more information, see *Scripting API for Spirent Avalanche Reference Manual*.

Troubleshooting Hardware and Software FAQs

How do I find error messages?

You can find software and hardware messages in the Event Log: Click the **Run** tab, and then the **Monitor** tab.

Commander saves the event log data for each test run. You can access the event log file (called `event.log`) by Analyzing Test Results in the Results tab.

What are common error messages found in the Event Log files and what do they mean?

- Receiving a parse error indicates incorrect URL syntax in your URL list.
- An active link error is caused when none of the test ports have an active link.
- A License key curbing error indicates the test configuration is exceeding the limits imposed by your current license key settings.
- A segmentation fault (segfault) might occur if the system can't handle resource allocation. If a segfault occurs, send [Spirent Communications Technical Support](#) a copy of the backtrace found in the Event Log for further troubleshooting. To retrieve the backtrace, select debug from the Event Log Level in the Run Configure tab. In debug mode, click the Event Log button in the Run Monitor tab to see the backtrace.
- A "throttling due to low memory" error message might occur if you have reached the limits of your memory pool. The appliance continues to run, however, it does not open new users or connections until the memory usage decreases.

For more information, see "Viewing Event Logs" and "Running a Test" in the Avalanche Commander Help.

How do I find system logs?

You can find various system logs in your installation logs directory (for example, `C:\Program Files\Spirent Communications\Avalanche\logs`). For each type of system log, a maximum of 10 files (10 MB each) is maintained, with the oldest (highest number) being deleted. (For example, `commander.log` becomes `commander.log.1`, `commander.log.2`, and so on, up to `commander.log.9`, then subsequently deleted.) There is no need for you to view these files during normal operation. All useful error messages appear in pop-up windows and the Event Log. Occasionally, [Spirent Communications Technical Support](#) may request these system logs to help troubleshoot the system.

What TCP stack does Avalanche use?

The Avalanche TCP stack is a custom stack written from scratch with many ideas derived from BSD-based stacks.

What is the DNS timeout on Avalanche?

Three seconds or three retries. You can set the number of DNS retries and interval between retries on the Client Profiles tab. For more information, see the Avalanche Commander Help.

What is the OS base?

QNX is the OS base on the Avalanche appliance.


Linux is the OS base for SmartBits chassis.

What is the Port Mapping for Avalanche?

Port numbering for Avalanche/Reflector 2500/2200 horizontal configuration:



Port numbering for Avalanche/Reflector 2500/2200 vertical configuration:



What is the difference between a simuser, connection, and transaction?

These specifications define load traffic as follows:

Connections. This specification denotes a TCP connection. It defines the number of simultaneous network connections initiated from Avalanche. This setting generates enough load to reach and sustain the desired number of open TCP connections. You can use this load specification with any TCP connection-based protocol including HTTP, FTP, and SMTP.

Open TCP connections depend on the connection establishment rate, system efficiency, and user behavior, such as Think Time. Specifying the same number of TCP connections on different systems does not mean the traffic generated will be the same. An open connection does not necessarily have nonstop activity.

Pay attention to the time it takes a TCP connection to complete, the time to SYN/ACK, and the number of successful versus failed transactions at a specified number of open TCP connections.

Connections/second or Connections/hour. This specification is often preferred by network equipment manufacturers for testing network devices. One TCP connection can contain up to hundreds of transactions, depending on client profile configuration and server response.

While the basic unit of load generated is SimUsers, Avalanche calculates this load type based on historical data. Avalanche dynamically adjusts the rate of user arrival (as in SimUsers/second or SimUsers/hour) to match the targeted connections/second or connections/hour rate.

SimUsers. This is an abbreviation for simulated or virtual users. Simuser applies to a user processing through an Action list one time. This load specification generates enough load to reach and maintain a target number of concurrent simulated users. It allows you to determine the maximum number of concurrent users your device, infrastructure, or system can handle. Use this specification if you want to keep applying load even after the device under test fails. The amount of traffic generated depends on the performance of the device under test. As the system slows down due to overloading, generally each user takes longer to process through the URL list, and the load "throttles back" and generates fewer new users.

SimUsers/second or SimUsers/hour. This specification maintains a target number of concurrent simulated users per second or hour, and provides a more realistic simulation of users generating traffic on a system. The load specification helps you identify where the system breaks down because the load may ramp up too quickly and will not "throttle back," but will continue independently of the behavior of the device tested.

Transactions. This specification defines the number of simultaneous transactions generated. This specification generates and maintains enough load to reach an outstanding number of active HTTP transactions, or GETs-in-progress. For example, HTTP 1.1 with Persistence allows you to execute multiple transactions in a single connection. Each transaction equates to the request and transfer of one object, which for a website is called a hit.

NOTE: The Transactions specification applies to HTTP and HTTPS transactions only. For certain protocols, such as FTP, DNS, Streaming, POP3, and SMTP, use the SimUsers or SimUsers/second.

Calibrating the load to generate a specific transaction/second while maintaining realistic traffic is a challenge. Since load is generated as new users arrive, existing users continue their transactions at non-uniform rates due to Think Time and multiple level-2 URLs (embedded objects). The resulting increase in transactions/second may not be smooth due to traffic bursts created by each individual user.

Transactions/second or Transactions/hour. This specification gradually ramps up the number of HTTP transactions either per second or hour for the duration of the test. Transactions/second or hour might not equal connections/second or hour when using HTTP 1.1 with Persistence.

NOTE: The Transactions specification applies to HTTP and HTTPS transactions only. For certain protocols, such as FTP, DNS, Streaming, POP3, and SMTP, use the SimUsers or SimUsers/second.

Calibrating the load to generate a specific transaction/second while maintaining realistic traffic is a challenge. Since load is generated as new users arrive, existing users continue their transactions at non-uniform rates due to Think Time and multiple level-2 URLs (embedded objects). The resulting increase in transactions/second may not be smooth due to traffic bursts created by each individual user.

What should I do if my test produces failures?

Click the **Trial Run** icon to create CSV and PCAP trace files. These result files help determine where the test is failing by displaying incomplete transactions. If your issues still exist, contact [Spirent Communications Technical Support](#).

What types of cables are needed for my hardware?

Use SC Connectors for 1801,1822, and 1832 models.

Use LC Connectors for 2200 and 2500 models.

The 220 units have two 10/100 Ethernet interfaces.

The 2500 units can have three different port configurations: All copper, all fiber, or two copper and two fiber.

When I run my test, why does the interface keep ARPing instead of sending traffic?

ARPing can occur if the interface does not have enough time to negotiate the link. There are two possible solutions to solve this issue: use at least a 60 second test start delay in Commander from the Run Configure tab, or use spanning tree within your testing network.

Why can't my test connect to the SmartBits chassis when I run my test?

This is most likely due to an IP mis-configuration on the chassis. Check the SmartBits chassis IP information to verify your settings are correct.

Why do I see resets at the end of my tests?

When the ramp down phase has expired, Avalanche Commander resets all outstanding connections, resulting in errors. When a test generates many errors at the end of the test, it is recommended that you modify your load specification to allow enough ramp-down time for all connections to complete successfully. Increase the ramp-down time to be larger than either the user session or the TCP timeout value (whichever is greater).

Why does Avalanche not respond in Commander, even though I'm able to Telnet into my hardware?

This is most likely caused by an application loading error from an incorrect OS version when you first upgraded your software.

Contact [Spirent Communications Technical Support](#) for more information.

Why should I use sample tests?

Avalanche ships with sample tests to better acquaint you with the product, and to help demonstrate its capabilities before you run your own tests. You generate sample tests using the Sample Test Wizard, then select the sample tests from the navigation tree in the Avalanche Commander window. For more information, see the Avalanche Commander Help.

How do I troubleshoot FTP problems such as throttling down, instability, and rejected connections?

Because FTP protocol has long data transfers, you may experience several of the following problems:

- Throttling down (out of memory)
- Instability
- Rejected connections; connection was unsuccessful/aborted.

These problems usually occur due to inaccurate configuration of the load profile. There are several parameters that can affect testing, such as load profile, Action list, memory, number of NICs, and NIC throughput. The key to successful testing is to have a balance between these parameters.

Consider an FTP test with the following parameters:

- 10,000 user sessions per second
- One user with 24 Actions
- Each Action retrieving 1kb file
- Only one 100baseT interface.

In this case, the test attempts to create $10,000 * 24 * 2$ or 480,000 FTP connections per second. Half of these connections (240,000 data connections) attempt to retrieve 1

kilobyte, which results in 240Mb per second. Additionally, the IP header traffic and FTP control connection traffic should be considered.

This example test results in constant out-of-memory conditions, instability, and almost all unsuccessful or aborted connections. To avoid these type of problems, it is recommended that you adhere to the following guidelines:

- Avoid using connections per second or sessions per second.
- The total load specified in the load profile per interface (in Mb) should not exceed the network throughput of that interface.
- Allow about 20% of reserved traffic for control, congestion, and so on.
- Do not create tests that exceed product specifications.

How do I troubleshoot unsuccessful streaming session?

If the streaming file duration is longer than the test duration, an unsuccessful session may occur at the end of the test. Therefore, your load specification should allow enough time for the streaming data to complete successfully for streaming tests.

Make sure that the streaming file that you request, exists on the streaming server (either Reflector or an actual streaming server). If you are running a test with Avalanche against Reflector, you must previously load the specified file as a content file. If you are running the test against an actual streaming server, you do not need to upload a content file.

Why don't I get better performance when I use more than one port with a test (assuming the ports are on the same card)?

You won't get better performance because the same processor controls all the ports on the card.

Why do my tests terminate prematurely when using SmartBits?

Excessive load on some SmartBits modules may cause tests to terminate prematurely. This is a memory issue, which may occur on a card configured for either the client or server. Additionally, if you set up a high-load test between client(s) and server(s) of significantly different performance levels (50% or more), tests may terminate prematurely. Once a test has terminated prematurely, you may have to power-cycle the client(s) and server(s) involved.

Why do my performance tests behave differently from version 5.2 to 7.0?

Performance tests that are converted from version 5.2 to 7.0 may not run exactly as expected, and could fail. Some changes in load patterns should be expected.

Why can other users access my appliance while my test is running, causing it to fail?

There is no mechanism to prevent other users from remotely accessing your appliance (for example, by using Telnet, FTP, and so on). Another user could change settings, perform a software upgrade, and reset your appliance, even if a test is running.

Why would my appliance become inaccessible through FTP and Telnet?

When the administration port IP address is updated through the GUI, the appliance becomes inaccessible through FTP and Telnet. To avoid FTP and Telnet lockups, update the matching DNS server when you update the administration port's IP address.

Why do I get RX Underrun/Overrun errors in my results summary CSV file?

Avalanche may drop packets when using Gigabit Ethernet interfaces, causing underrun and overrun errors, noticeable in the results summary CSV file. These errors are due to the current Avalanche architecture. They are harmless, and can safely be ignored.

What do I use for my installation password?

Send an e-mail message to [Spirent Communications Technical Support](#) to get the installation password.

Contacting Spirent Communications

To obtain technical support for any Spirent Communications product, please contact our Support Services Department using any of the following methods:

U.S.

E-mail: support@spirentcom.com

Web: <http://support.spirentcom.com>

Toll Free: +1-800-SPIRENT (+1-800-774-7368) (U.S. and Canada)

Phone: +1 818-676-2616

Fax: +1 818-880-9154

Operating Hours: Monday through Friday, 06:00 to 18:00 Pacific Time

Europe

E-mail: support.europe@spirentcom.com

Web: <http://support.spirentcom.com>

Phone: +33 (0) 1 61 37 22 70

Fax: +33 (0) 1 61 37 22 51

Operating Hours: Monday through Thursday, 09:00 to 18:00,

Friday 09:00 to 17:00, Paris Time

China

E-mail: supportchina@spirentcom.com

Web: <http://support.spirentcom.com.cn>

Toll Free: +1-800-810-9529 (mainland China only)

Phone: +86 10 8233 0033 (rest of the world)

Fax: +86 10 8233 0022

Operating Hours: Monday through Friday, 09:00 to 18:00 Beijing Time

The latest versions of user manuals, application notes, and software and firmware updates are available on the Spirent Communications Customer Service Center websites at <http://support.spirentcom.com> and <http://support.spirentcom.com.cn> (China).

Information about Spirent Communications and its products and services can be found on the main company websites at www.spirentcom.com and www.spirentcom.com.cn (China).

Company Address

Spirent Communications, Inc.

26750 Agoura Road

Calabasas, CA 91302

USA

All of the company names and/or brand names and/or product names referred to in this document, in particular, the name "Spirent" and its logo device, are either registered trademarks or trademarks of Spirent plc and its subsidiaries, pending registration in accordance with relevant national laws. All other registered trademarks or trademarks are the property of their respective owners.

2005 Spirent Communications, Inc. All Rights Reserved.