

2.6 - Scalability and Performance Test Methodology

Scalability and Performance Test Methodology — Appvance IQ documentation.

- [Scalability and Performance Test Methodology](#)
- [Applying the Method](#)
- [Background and Goals](#)
- [Add Calibration Testing to the Performance Test Plan](#)
- [Calibration What-if Chart](#)
- [Identifying Service Performance Metrics](#)
- [Identify Test Environment Hardware and Software](#)
- [Method for Black Box and White Box Tests](#)
- [Scalability Index](#)
- [Scalability and Performance Kit Contents](#)
- [Appvance Method](#)
- [Understanding the Tests](#)
- [Understand TPS in AIQ Performance Testing](#)
- [User Goal Oriented Testing \(UGOT\)](#)
- [Performance Testing Primer for AIQ](#)
- [Test Designer Performance Scalability: Maximizing the number of concurrent browsers per node](#)

Scalability and Performance Test Methodology

Use this article to plan AIQ performance tests that define goals, load model, environment readiness, execution controls, and result interpretation before a test is run.

What changed in this article

- Kept the title but refreshed the methodology around current AIQ performance testing practices.
- Added checks for business objective, workload model, test data, environment parity, monitoring, ramp-up, duration, success criteria, and result review.
- Added guidance to separate application bottlenecks from AIQ/Test Node capacity constraints.

Validation workflow

1. Define the performance question, workload model, success criteria, and monitoring plan before building the test.
2. Validate data, environment readiness, Test Node capacity, and baseline behavior with a small run.
3. Run the planned test and review throughput, response time, errors, saturation, and evidence before making conclusions.

Review note: This page now reads as a practical methodology checklist instead of a generic performance article.

Last updated: October 2, 2026

If you have feedback on this topic, contact documentation@appvance.ai.

Applying the Method

The Appvance method evaluates scalability, performance, and developer productivity across different implementation approaches for Internet application services. The methodology compares conventional application-stack implementations with native XML and data-processing technologies, focusing on both runtime performance and development efficiency.

Scalability and performance are evaluated from two primary perspectives:

1. **Binding and data-processing acceleration**

Evaluates different approaches for processing XML-based data, including object binding, streaming parsers, DOM-based processing, and XQuery. The use cases compare execution performance, resource consumption, scalability, and developer productivity across the different approaches.

2. **Mid-tier caching and service acceleration**

Evaluates the impact of caching and data-storage strategies on service performance. The use cases compare native XML databases (XML DB) and relational databases (RDBMS) across different XML message sizes and database operations, including insert, update, delete, and query operations.

Background and Goals

Software architects and developers choose XML parsing techniques, service libraries, encoding techniques, and protocols when building services using SOA techniques. Each choice has an impact on the scalability and performance of the finished service. In this example we have 3 goals:

1. Explain the changing landscape of APIs, libraries, encoding techniques, and protocols to software architects and developers. The current generation of technology choices changes approximately every 6 to 9 months. For instance, JAXB 1.0 is replaced by JAXB 2.0 and WebLogic Server 8.1 is replaced by WebLogic Server 9
2. Identify and use real-world, test use case scenarios showing software architects and developers how to choose technology based on their service goals
3. Deliver code compatible with the current techniques for building functional and scalability tests (black-box, unit, agile test-first). Many vendor tutorials are not compatible with what was used to build Appvance. While these packages could be modified with a lot of work, a better resolution was to develop kits with commercial, incorporating public information accessible to anyone

By pursuing these techniques we build a reusable method for evaluating Internet application performance and system scalability, plus the results feed basic business needs of cost/benefit and feature/function analysis including:

- Reduced hardware costs and reduced per-CPU licensed software by running more efficient services built with more efficient tools.
- Increased efficiency with lower network and processor bandwidth.
- Reduced time-solving interoperability problems with more effective tools.

The kit arms business managers and software developers with the evidence needed to recommend and adopt new solutions internally, and get their projects funded.

Last updated: February 22, 2023

If you have any feedback on this topic, please contact us at documentation@appvance.ai

Add Calibration Testing to the Performance Test Plan

Recommended guidance

Add calibration testing to the performance test plan before running large scalability tests. Calibration confirms that the script, data, environment, load model, monitoring, and result collection are ready for higher-volume execution.

Calibration goals

- Validate that the scenario works reliably with a small user count.
- Confirm test data, DPLs, accounts, and environment access are sufficient.
- Measure baseline response times, errors, throughput, and resource use.
- Identify bottlenecks in the test script, Test Nodes, network, or application under test.
- Confirm monitoring, logs, reports, and dashboards capture the required evidence.

Planning checklist

1. Run a single-user validation before adding concurrency.
2. Increase load in controlled steps and review errors at each level.
3. Separate application bottlenecks from load-generator or environment limits.
4. Document calibration results and use them to refine the final test plan.

Last updated: October 5, 2026

Calibration What-if Chart

Calibration testing helps identify how a system behaves as the load level changes and provides a basis for selecting appropriate test levels. However, a calibration test cannot account for every possible system behavior or bottleneck.

The **Calibration What-If Chart** provides a set of common test observations, their likely causes, and recommended next steps. Use these observations to determine whether additional calibration is required before proceeding with the full performance test.

Test Experience	Likely Problem	What to Do Next
Increase in CRs with a decrease in TPS	Check the average response time.	Run a test case comparing response times as CRs increase. Identify the maximum acceptable response time and use it to determine the appropriate test level.
Increase in CRs with an increase in TPS	The test may not be calibrated for sufficiently high CRs or payload sizes.	Run additional calibration tests to determine the appropriate Scalability Index and corresponding CR and payload-size levels.
Increase in CRs with little change in TPS	CR levels may already be too high. Check CPU utilization if performing a stateless test.	Run another test with the CR level reduced, for example, by 50%, and compare the results.
Server CPU reaches approximately 95% utilization and increasing CR levels increases TPS	The test is likely exercising a stateful system behavior.	Run a test case to determine the service-under-test Scalability Index.
Consumer CPU reaches approximately 95% utilization and increasing CR levels decreases TPS	CR levels may be too high for the number of load-generating consumers.	Add more load-generating consumers and repeat the calibration test.
Consumer CPU utilization is approximately 15%, server CPU utilization is approximately 30%, and increasing CR levels produces little change in TPS	The network may be saturated or otherwise limiting throughput.	Check network bandwidth utilization. Consider adding network adapters to the server or using a faster network configuration.

Calibration testing provides the information needed to identify practical TPS levels for the system under test based on its current hardware, software, network, and workload configuration.

After calibration, use the resulting test levels to run the **full performance test** and evaluate the system under the intended workload.

Note: The observations in this chart are guidelines rather than definitive diagnoses. Actual bottlenecks may vary depending on the application architecture, workload, payload size, network configuration, and test environment.

Last updated: August 14, 2026

If you have any feedback on this topic, please contact us at documentation@appvance.ai.

Identifying Service Performance Metrics

Identifying Service Performance Metrics

Overview

Scalability and performance testing provides meaningful business value when it is performed using a consistent, repeatable, and measurable methodology.

A performance test should do more than determine whether a service is "fast enough." It should provide actionable information about how the service behaves under different workloads, how much traffic it can support, and where performance or scalability begins to degrade.

A complete performance testing methodology should address three fundamental areas:

1. **Selecting representative test cases**
2. **Executing accurate and repeatable test runs**
3. **Analyzing the results to identify scalability, capacity, and performance characteristics**

These areas work together to provide a clear understanding of how a service behaves under realistic and increasing workloads.

1. Select the Right Test Cases

The first step in performance testing is determining what should be tested.

Not all services behave the same way, and therefore not all services should be tested using the same workload.

For example, a service that handles a high volume of small requests may require a different test strategy from a service that receives large messages periodically.

Consider the following examples:

High-volume service

Small requests
High request frequency
Many concurrent users

The primary concern may be:

- Maximum throughput.
- Transactions per second (TPS).
- Concurrent request capacity.
- Response time under sustained load.

Large-message service

Large request payloads
Lower request frequency
Higher processing cost per request

The primary concerns may instead include:

- Response time.
- Network throughput.
- Processing time.
- Memory consumption.
- Maximum supported message size.

Multiple-interface service

A service may expose several interfaces or API operations, each with different performance characteristics.

For example:

```
POST /login
GET  /products
POST /cart
POST /checkout
```

Testing only one endpoint may not accurately represent the workload generated by a real user.

The selected test cases should therefore represent the **actual business and user workflows** that the service is expected to support.

2. Define Representative Workloads

After identifying the relevant test cases, define the workload that will be applied to each test.

Important workload dimensions include:

- Number of concurrent users or requests.
- Request frequency.
- Message or payload size.
- Message complexity.
- Transaction mix.
- Duration of the test.
- Ramp-up and ramp-down behavior.
- Environment and configuration.

For example, a test plan could include:

Scenario	Concurrent Users	Payload	Duration
Baseline	10	Small	10 min
Normal Load	100	Small	30 min
High Load	500	Small	30 min
Large Payload	100	Large	30 min
Stress	1,000	Small	Until degradation

The objective is not simply to generate as much traffic as possible.

The objective is to understand **how the service behaves as the workload changes**.

3. Capture Real Application Traffic

Whenever possible, performance scenarios should be based on realistic application behavior.

For browser-based workflows, **Test Designer Instant Replay** can be used to capture the network activity generated during recording or playback.

A typical workflow is:

```
Test Designer
  |
  v
Browser UX workflow
  |
  v
Instant Replay
  |
  v
Captured network traffic
  |
+----> HAR
+----> CSV
+----> JS
```

The captured traffic can help identify the API requests associated with a real user workflow.

For example:

```
Login
  |
  v
Search Product
  |
  v
Add Product
```

|
v
Checkout

The corresponding API traffic can then be reviewed and, when appropriate, reused to build protocol-level performance scenarios.

This approach helps maintain a connection between **real user behavior** and the workload used during performance testing.

4. Execute Accurate and Repeatable Test Runs

A single performance test run is rarely sufficient to understand scalability.

Performance results can be affected by many factors, including:

- Current system load.
- Network conditions.
- Test environment configuration.
- Database activity.
- Application state.
- Test data.
- Infrastructure resources.
- Number of concurrent requests.

For this reason, performance testing should consist of multiple controlled and repeatable executions.

A typical progression may look like:

```
Run 1  → Baseline
Run 2  → Normal load
Run 3  → Increased load
Run 4  → High load
Run 5  → Stress
Run 6  → Repeat high load
```

Repeated executions make it easier to distinguish consistent system behavior from temporary variations.

Manual or ad-hoc recording of results is not sufficient for large-scale performance testing. Test execution and data collection should be automated whenever possible.

Automation provides:

- Consistent execution.
 - Repeatable workloads.
 - Standardized measurements.
 - Automated result collection.
 - Easier comparison between runs.
 - Reduced human error.
-

5. Define the Performance Metrics

A performance test is only useful if the results are measured using clearly defined metrics.

The following metrics are commonly used to evaluate service performance.

Transactions Per Second (TPS)

TPS measures the number of transactions processed by the service per second.

For example:

TPS = 1,500

means that the test generated approximately 1,500 transactions per second.

TPS is particularly useful when evaluating service throughput and capacity.

Response Time

Response time measures how long it takes for a request to complete.

For example:

Average Response Time = 250 ms

Average response time is useful, but it should not be the only response-time metric considered.

Averages can hide slow requests.

Percentile Response Times

Percentiles provide a better understanding of the distribution of response times.

Common metrics include:

- **p50** — 50% of requests completed within this time.
- **p90** — 90% of requests completed within this time.
- **p95** — 95% of requests completed within this time.
- **p99** — 99% of requests completed within this time.

For example:

p50 = 200 ms
p90 = 350 ms
p95 = 500 ms
p99 = 1,200 ms

This tells us that while most requests are relatively fast, a small percentage are taking significantly longer.

Error Rate

Error rate measures the percentage of requests that fail.

For example:

Total Requests = 1,000,000
Failed Requests = 2,000

Error Rate = 0.2%

Errors should be analyzed alongside throughput and response time.

A system that processes 5,000 TPS but returns a high percentage of errors cannot be considered healthy.

Concurrent Users or Requests

Concurrency represents how many users or requests are active at the same time.

Increasing concurrency is one of the primary ways to evaluate scalability.

For example:

100 users	→ 1,000 TPS
500 users	→ 3,500 TPS
1,000 users	→ 5,000 TPS
2,000 users	→ 5,100 TPS

In this example, increasing concurrency beyond 1,000 users produces very little additional throughput.

That may indicate that the system is approaching a capacity limit.

Resource Utilization

Performance testing should also consider the resources supporting the service.

Depending on the architecture, relevant measurements may include:

- CPU utilization.
- Memory utilization.
- Disk utilization.
- Network throughput.
- Database utilization.
- Connection pools.
- Thread pools.
- Queue depth.

These metrics help explain **why** performance changes as load increases.

6. Understand Scalability

Scalability describes how the system responds when the workload increases.

Consider this example:

Concurrent Users	TPS
10	500
50	1,200
100	2,000
500	3,500
1,000	4,000
2,000	4,050

Initially, increasing concurrency produces a significant increase in throughput.

However, after approximately 1,000 users, TPS begins to level off.

This indicates that the service may be approaching its practical capacity.

At the same time, response times may begin to increase:

Users	TPS	p95 Response Time
100	2,000	300 ms
500	3,500	450 ms
1,000	4,000	800 ms
2,000	4,050	2,500 ms

Although TPS continues to increase slightly, response time has degraded significantly.

This is an important scalability finding.


7. Analyze the Relationship Between Workload and

Performance

The objective of performance testing is not simply to collect numbers.

The objective is to understand the relationship between workload and system behavior.

Important relationships include:



For example, increasing concurrency may initially improve throughput.

Eventually, the system may reach a saturation point where:

- TPS stops increasing.
- Response times increase.
- Error rates increase.
- CPU or memory utilization becomes high.
- Database or network resources become constrained.

This point is often more important than the maximum TPS observed during the test.

8. Identify the Service Capacity

Performance testing should help establish the practical capacity of the service.

For example:

Maximum observed throughput: 4,100 TPS

But at 4,100 TPS:

p95 response time = 3 seconds

Error rate = 4%

It would be misleading to state that the service supports 4,100 TPS as a healthy operating point.

Instead, the team might determine that:

Recommended operating capacity:

3,500 TPS

p95 response time:

< 500 ms

Error rate:

< 1%

This provides a much more useful capacity target.

9. Compare Different Environments and Configurations

Performance testing should also be performed across relevant environments and configurations.

For example:

```
QA
|
+---- Configuration A
|
+---- Configuration B
```

UAT

|

+---- Configuration A

|

+---- Configuration B

Production-like

|

+---- Configuration A

|

+---- Configuration B

The same workload can then be used to determine whether differences in configuration or infrastructure affect:

- TPS.
- Response time.
- Error rate.
- Resource utilization.
- Maximum capacity.

For comparisons to be meaningful, the test conditions should remain as consistent as possible.

10. Establish a Baseline

Before evaluating whether a change improves or degrades performance, establish a baseline.

A baseline should capture the expected behavior of the service under a known workload.

For example:

Baseline

Concurrency: 500

TPS: 3,500

p95: 450 ms

Error Rate: 0.2%

CPU: 65%

After an application or infrastructure change, the same workload can be executed again.

New Version

Concurrency: 500

TPS: 3,200

p95: 700 ms

Error Rate: 0.8%

CPU: 82%

The comparison immediately shows that the new version performs worse under the same workload.

This type of baseline comparison is much more useful than simply saying that a test "passed."

11. Turn Results into Actionable Knowledge

The final step is to interpret the results and produce conclusions that can support technical or business decisions.

A useful performance conclusion should answer questions such as:

- What workload can the service currently support?
- What is the maximum sustainable throughput?
- At what point does response time degrade?
- When does the error rate become unacceptable?
- Which resource becomes the bottleneck?
- Does the service scale as concurrency increases?
- How does payload size affect performance?
- How does message complexity affect performance?
- Does the service behave consistently across environments?
- Did a new application version improve or degrade performance?

For example:

“ The service scales from 100 to approximately 500 concurrent users, increasing throughput from 1,800 to 3,500 TPS. Beyond 500 users, throughput increases only marginally while p95 response time increases significantly. CPU utilization reaches 90% at approximately 1,000 concurrent users, indicating CPU saturation

as the primary scalability constraint.

This is actionable information because it provides a basis for deciding whether to optimize the application, increase infrastructure capacity, or establish a lower operating limit.

Performance Testing Methodology

The complete methodology can be summarized as follows:

1. Identify user workflows
|
v
2. Select representative test cases
|
v
3. Define workload characteristics
|
v
4. Capture realistic API/network traffic
|
v
5. Create repeatable performance scenarios
|
v
6. Execute multiple controlled test runs
|
v
7. Collect performance and resource metrics
|
v
8. Analyze throughput, latency, errors,
concurrency, and resource utilization
|
v

9. Identify scalability and capacity limits

|

v

10. Compare against baselines

|

v

11. Produce actionable conclusions

Appvance Performance Testing Approach

The Appvance approach combines realistic user workflows with automated performance testing and measurable service-level results.

A typical workflow can begin with **Test Designer**, where the user journey is recorded or executed.

Instant Replay can then capture the network traffic generated by that workflow.

The captured traffic can be reviewed or exported and used as a foundation for protocol-level testing.

This provides a connection between:

User Experience

|

v

Test Designer

|

v

Network/API Traffic

|

v

Instant Replay

|

v

Protocol-Level Test

|

v

Performance Test

|

v

Metrics and Analysis

|

v

Scalability / Capacity Findings

This approach allows teams to evaluate both the **user-facing workflow** and the **service behavior behind that workflow**.

Key Takeaways

Effective scalability and performance testing requires more than generating a large number of requests.

A successful methodology should:

1. **Test realistic user and service workflows.**
2. **Use representative workloads and data.**
3. **Automate and repeat test execution.**
4. **Measure consistent and meaningful metrics.**
5. **Evaluate throughput as workload increases.**
6. **Monitor response-time percentiles and error rates.**
7. **Monitor infrastructure and service resources.**
8. **Identify capacity and saturation points.**
9. **Compare results against established baselines.**
10. **Turn test results into actionable technical and business conclusions.**

The ultimate goal is not simply to determine whether a service can process a particular number of requests.

The goal is to understand **how the service behaves as demand increases, where its limits are, and what actions are required to maintain acceptable performance and scalability**.

Last updated: August 12, 2026

For feedback or questions about this documentation, contact the appropriate Appvance documentation or support team.

Identify Test Environment Hardware and Software

Recommended guidance

Document the complete test environment before running scalability or performance tests. Results are only useful when the hardware, software, network, data, monitoring, and AIQ execution configuration are known and repeatable.

Capture these details

- Application servers, databases, caches, queues, load balancers, and external dependencies.
- CPU, memory, disk, network, operating system, runtime, browser, and driver versions.
- AIQ version, Test Node count, node sizing, browser mode, and execution geography.
- Data volume, account pool, DPL/CSV sources, and cleanup/reset process.
- Monitoring tools, metrics, log locations, and dashboard links.

Why it matters

Environment changes can explain different response times, error rates, and scalability limits. Record changes between calibration, baseline, stress, and final runs so results can be compared fairly.

Last updated: October 5, 2026

Method for Black Box and White Box Tests

Testing services for scalability and performance differs from testing individual software components or application code. Service performance testing focuses on understanding how a service behaves as the workload increases, including changes in concurrent requests, message sizes, message complexity, and response handling.

The Appvance methodology combines **black-box** and **white-box** testing. Black-box testing establishes how the service performs from the perspective of its consumer, while white-box techniques can then be used to identify the internal operations responsible for performance bottlenecks.

The recommended approach is:

- 1. Establish a baseline using black-box performance testing.**
Create a baseline performance metric, or **Scalability Index**, by measuring Transaction Per Second (TPS) at the service consumer. Test the service using different message sizes, message schema complexities, and levels of concurrent requests.
- 2. Calibrate and compare the test environment.**
When comparing multiple servers, consumers, brokers, or other service configurations, identify the **Performance Index** for each configuration. Normalize the test parameters and environment so that differences in infrastructure or workload do not produce misleading performance results. This calibration ensures that subsequent comparisons are meaningful.
- 3. Identify internal performance bottlenecks.**
After establishing the Performance Index, use white-box techniques such as code profiling to identify the operations that consume the most processing time when handling requests. Use this information to determine where optimization can provide the greatest performance improvement.
- 4. Optimize and repeat the measurements.**
Apply the identified optimizations and repeat the performance tests. Compare the new Performance Index with the previous baseline to determine whether the changes improved scalability or performance.
- 5. Validate the final performance.**
Run the Performance Index again under the calibrated test conditions and analyze the results. The final measurements provide an objective basis for evaluating the effectiveness of the optimizations and the scalability of the service.

This combination of black-box performance measurements and targeted white-box analysis provides the basis for the Appvance methodology for testing services for scalability and

performance.


Last updated: August 12, 2026

If you have feedback about this topic, please contact documentation@appvance.ai.

Scalability Index

The **Scalability Index** describes how a service responds to increasing concurrent request (CR) levels and payload sizes. It is used during calibration to identify the range in which the service operates efficiently and to determine appropriate load levels for the Full Test.

The chart below illustrates three representative regions of a scalability curve:



- **View A — Underdriving:** Increasing the CR level results in a corresponding increase in TPS. The service is still able to process additional load efficiently, indicating that the current load level may be below the service's optimal operating range.
- **View B — Optimal operating range:** The service reaches its most useful operating range. Increasing the CR level no longer produces proportional throughput gains, but the service continues to operate at an acceptable TPS level. This range is typically used to define the CR levels for the Full Test.
- **View C — Overdriving:** Increasing the CR level produces little or no additional TPS. The service may be saturated, or another system resource may have become the limiting factor. Increasing the load further is therefore unlikely to provide useful scalability information.

The goal of calibration is to identify the **View B range** rather than simply maximize the CR level. Testing at too few requests may underutilize the service, while testing at excessive levels may introduce saturation that obscures the service's useful scalability characteristics.

Once the appropriate CR range has been identified, use those levels in the **Full Test**. The same calibrated levels can also be used when comparing alternative APIs, implementations, or products so that their performance results can be evaluated under equivalent workloads.

Identifying Bottlenecks

Resource utilization can help determine what limits scalability during calibration. Depending on the type of service and test, review metrics such as:

- CPU utilization
- Memory utilization
- Network bandwidth
- Load-generator or consumer utilization
- Service response time and TPS

The relevance of each metric depends on the architecture and behavior of the service. For example, CPU and memory utilization can provide useful information for stateless services, while resource utilization may be less representative of scalability for stateful services where other factors determine the system's behavior.

The **Scalability Index should therefore be considered together with system-resource metrics and response-time behavior**, rather than relying on TPS alone.

Last updated: August 13, 2026

If you have any feedback on this topic, please contact us at documentation@appvance.ai.

Scalability and Performance Kit Contents

Recommended guidance

A scalability and performance kit should contain the information needed to plan, run, analyze, and repeat an AIQ performance test. Keep the kit organized so another reviewer can understand the load model, environment, scripts, data, results, and conclusions.

Recommended contents

- Test objectives, scope, assumptions, and success criteria.
- Load model, user journeys, concurrency, ramp-up, duration, and pacing.
- Scenario Editor or service assets, DPLs, CSV files, and dependency notes.
- Environment inventory, AIQ/Test Node configuration, and monitoring setup.
- Calibration, baseline, stress, endurance, and final result summaries.
- Defects, risks, tuning actions, retest notes, and final recommendations.

Quality checks

1. Remove duplicate or outdated artifacts before sharing the kit.
2. Mask credentials, tokens, customer data, and sensitive URLs where required.
3. Link to AIQ reports and dashboards instead of copying unsupported screenshots only.
4. Include enough context for the test to be repeated later.

Last updated: October 5, 2026

Appvance Method

Phase	Goal	Checkpoint
Planning	Answer the question: How will this test benefit the organization?	Write Test Plan document.
Definition	Identify use cases, test cases and scenario, and test environment (hardware, software and network)	Add the use cases, test cases and test scenario to the Test Plan.
Calibration Test	Calibrate the test cases to the test environment.	Identify the use cases driving the test environment to its maximum throughput (as measured in TPS from the client).
Optimize	Modify the service and / or test environment to optimize for the best performance based on what was learned in the Calibration Test.	Amend the Test Plan to add the changes for optimization.
Full Test	Run the <i>TestScenario</i>	Successful run of <i>TestScenario</i> .
Results Analysis	Identify test result metrics and trends against the <i>TestScenario</i> goals.	Present results and achieve adoption by management.

The **Appvance Test Methodology** uses three key concepts to define and organize tests:

- **Use Case** – describes the functionality or approach being tested. For example, a test comparing XML parsing techniques could have two use cases: one using the Xerces DOM parser and another using JAXB XML binding.
- **Test Case** – defines the specific inputs or conditions under which an individual test is executed. For example, when comparing service throughput for different message payload sizes, two test cases could be defined: one using a 500-byte XML document and another using a 10,000-byte XML document.
- **Test Scenario** – defines the complete execution by combining the required use cases and test cases into a single test.

The Appvance methodology differs from traditional software testing in an important way: **it is designed to evaluate not only whether functionality works correctly, but also how the system behaves under different conditions and workloads.**

This distinction is particularly important for scalability and performance testing, where the same functionality may need to be evaluated across different payload sizes, concurrency levels, environments, and execution conditions.



Here is a brief explanation of the terminology used in [The Appvance Test Methodology](#)

Last updated: August 17, 2026

If you have any feedback on this topic, please contact us at documentation@appvance.ai

Understanding the Tests

The Appvance Test Methodology defines common test goals used when testing Internet application services. Because application architectures and service technologies continue to evolve, the methodology is not limited to these test goals.

The following concepts are used to understand the tests defined in the methodology:

- **Test Name** identifies the test being performed.
- **Test Benefit** describes the primary benefit or objective of the test.
- **Test Type** determines whether the test is **stateful** or **stateless**. Each type requires a different testing strategy and addresses different performance goals.
- **Scalability Index Parameters** define the test inputs used to measure the Scalability Index. These inputs represent expected usage patterns and help determine how service performance is expected to behave under production-like conditions.

Test Name	Test Benefit	Test Type	Parameters Related to the Scalability Index
Service Interface	Decrease time for service request responses to lower network bandwidth and server hardware costs.	Stateless	Message size and concurrent request level.
XML Parsing	Decrease time for routing service messages to lower network bandwidth and server hardware costs.	Stateless	Schema complexity (depth and element count), document size and concurrent request level.
Data Persistence	Decrease time for storing and retrieving messages to lower network bandwidth, server hardware and disk costs.	Stateful	Schema complexity (depth and element count), document size, concurrent request level.
Data Transformation	Decrease time for transforming a message into a given XML schema to lower network bandwidth and server hardware costs.	Stateless	Source and destination schema complexity (depth and element count), request and output document size, concurrent request level.
Data Aggregation and Federation	Decrease time for responding to service requests requiring up-stream data to reduce network bandwidth and server hardware costs.	Stateful	Schema complexity (depth and element count) for up-stream services, data persistence quantities, message time-to-live (TTL) values.
Data Mitigation	Reduce time when a service is unavailable at peak usage to improve service availability and user satisfaction.	Stateful	Schema complexity (depth and element count) for each request, document size, concurrent request level.

Test Types

Stateless Testing

Stateless testing is used for services that process each request independently of other requests. The test evaluates how the service responds as factors such as **concurrent requests** and **message payload size** change.

Stateful Testing

Stateful testing is used for services that maintain state or depend on information from previous requests. This can include services that provide:

- Data persistence
- Workflow or transaction processing
- Message queuing
- User sessions
- Data indexing

Stateful testing evaluates how these state-dependent operations affect service performance and scalability under different workloads.

Scalability Index Parameters

The parameters associated with the **Scalability Index** define the inputs used during the test, such as concurrency levels and payload sizes.

These parameters represent the expected usage pattern of the service in production and are used to determine how service performance changes as the workload increases.

The Appvance Test Methodology defines the series of steps required to plan, execute, and analyze these tests.

For more information, see [Appvance Method](#)

Last updated: August 19, 2026

If you have any feedback on this topic, please contact us at documentation@appvance.ai

Understand TPS in AIQ Performance Testing

Recommended guidance

Transactions per second (TPS) measures how many completed business transactions or requests are processed during a test window. In AIQ performance testing, define TPS carefully before comparing runs because different teams may count transactions, pages, API calls, or user journeys differently.

What to define

- The exact transaction boundary: request, step, business action, or full journey.
- The measurement window used for average, peak, and sustained TPS.
- Whether failed, timed-out, retried, or partially completed transactions are included.
- The user load, ramp-up, pacing, test data, and environment used for the result.
- The related metrics: response time, error rate, throughput, CPU, memory, and network utilization.

How to use TPS

1. Validate the script and data at low load before treating TPS as meaningful.
2. Compare TPS only between runs with the same scenario, load model, and environment.
3. Watch for increasing error rate or response time when TPS rises.
4. Document whether the reported TPS is target, achieved, peak, or sustained.


Last updated: October 6, 2026

User Goal Oriented Testing (UGOT)

Countering software developers by saying "If you design for every possible user, no individual user will have their goals met when they use your software application!" Coverage tests are typically pointless, as users typically take a path through the functions in service, one after another, like chain links. The agile development community approaches this problem by recommending a test-first strategy, which is to write a unit test of a class before writing the class itself. At build-time, the compiler environment compiles the object code for the class and then runs the unit test against the compiled code. The unit test completes successfully by receiving example data and validating the response. If the class returns an invalid response, the unit test throws an exception that the build and deploy environment handles.

Unit testing and agile development methods help but are not a complete solution to UGOT techniques. For instance, a test first is usually only carried out on a unit level. SOA deploys applications as a collection of services, so testing individual units miss most of the big problems that occur during SOA integration and deployment. UGOT modeled tests check service as an individual user would – by picking one feature after the next in a chain of service requests, which is a more global and inclusive testing mechanism.

UGOT is ideal to understand performance and scalability. Internet applications must deliver real value to those who build and operate services. Understanding what is being tested is the heart of the issue, as can be seen in the following chart that shows the building blocks of Java development tools for building applications.




The components build on each other in three tiers:

- At the bottom are the fundamental components for soap bindings, XML parsing, JAVA inter-application messaging services (JMS), and clustering

- Building on these is service bus components for services to inter-operate at a message level
- The top tier provides interoperability at the application level

Given these building blocks and tiers, testing should be done across the tiers, where key performance bottlenecks in soap Bindings and JMS services may happen. Each of the components in the above chart impacts the scalability and performance of application implementation. Additionally, understanding the testing goal can be seen by reviewing the definition of service architecture as shown in the following illustration.



Which should be tested and when should it be tested? Performance tests normally check at least two. For instance, one test may check a consumer and service, and another check a consumer and a broker.

The Appvance Methodology lists options to understand what part of the system would benefit most from scalability and performance analysis.

Last updated: June 02, 2023

If you have any feedback on this topic, please contact us at documentation@appvance.ai

Performance Testing Primer for AIQ

Performance Testing Primer for AIQ

Last reviewed: September 21, 2026

This primer explains the baseline concepts for planning performance and scalability tests in AIQ. Use it before designing load profiles or interpreting reports.

When to use this guidance

- You are preparing a new performance test effort.
- Stakeholders need shared terminology for load, throughput, latency, and scalability.
- You need to define success criteria before running tests.

Recommended approach

1. Define the business workflow, population, and target volume.
2. Establish baseline response times and throughput with a small load.
3. Create realistic ramp-up, steady-state, and ramp-down patterns.
4. Monitor application, database, network, and AIQ test node capacity during execution.
5. Compare results against agreed thresholds and document bottlenecks.

Note: A performance test without clear objectives usually produces unclear results. Define goals before execution.

Last updated: September 21, 2026

If you have feedback on this topic, contact documentation@appvance.ai.

Test Designer Performance Scalability: Maximizing the number of concurrent browsers per node

For performance testing, you can maximize the number of concurrent browsers that can be opened in a node.

For AWS EC2 nodes running AIQ, the number of concurrent Chrome browsers that can be assigned per node is around 20 to 30, depending on the application.

The main bottleneck to achieving more concurrent browsers is the speed of the storage volume as measured in IOPS (input/output operations per second). A typical 60GB EBS volume is not rated higher than 3000 IOPS.

To avoid the bottleneck of volume speed, you can leverage **AWS EC2 instances that include NVME SSD volumes**. These types of volumes are very fast, but they are ephemeral storage volumes. Any data stored in these volumes will be lost once the instance is stopped or terminated. Also, the data in the NVME SSD volumes are not saved in the AMI images.

So the process of starting appvance in one of this nodes involves several initial steps which are coded in the provided **nvme_setup.sh** script recommended for performance nodes. The script does:

1. Mount the NVME SSD storage as a device visible to the OS
2. Copy GraalVM and aiq_distribution folder to the NVME SSD
3. Modify global.json in aiq_distribution/NetworkProxy/config/global.json so Chrome browser has the following option added:

```
{
  "name": "experimentalDS3",
  "value": "--disk-cache-dir=cache|--disk-cache-size=1"
}
```

4. Start AIQ from the NVME SSD drive.

Setting up a Amazon Linux 2 node for performance testing

1. Add the following file (nvme_setup.sh) to the aiq_distribution folder in the Amazon Linux 2 instance.

This script will mount the NVME SSD volume, if found, and then copies the Graalvm and aiq_distribution folders to the SSD drive.

Note: This script assumes the AIQ is located at /opt/aiq/aiq_distribution. Edit the script if the AIQ location is different.

nvme_setup.sh

```
#!/bin/bash
# This script assumes:
# - using a AWS ec2 instance with an ephemeral SSD NVME volume (for example m5ad,
#   c5d instance families)
# - AIQ is installed at /opt/aiq/aiq_distribution, modify below if not.
# - Linux has NVME volume drivers (included in Amazon Linux 2 OS)
# - graalvm is installed and $JAVA_HOME is set to its location
if [ -e '/dev/nvme1n1' ] ; then
    echo "Mounting NVME volume as /data"
    sudo file -s /dev/nvme1n1
    sudo mkfs -t xfs /dev/nvme1n1
    if [ ! -d '/data' ] ; then
        sudo mkdir /data
    fi
    sudo mount /dev/nvme1n1 /data
    sudo chown ec2-user:ec2-user /data
    sleep 5
    FOLDER=$(basename $JAVA_HOME)
    sudo cp -R $JAVA_HOME /data
    export JAVA_HOME=/data/${FOLDER}
    export PATH=/data/${FOLDER}/bin:$PATH
    echo "JAVA_HOME: $JAVA_HOME"
    sudo cp -R /opt/aiq/aiq_distribution /data/aiq_distribution
    sudo chown -R $USER:$USER /data
    cd /data/aiq_distribution/TestNode
else
    echo "Skipping mounting /data as instance do not have an NVME volume
(/dev/nvme1n1)"
```

```
fi
```

```
bash startup.sh
```

2. Modify global.json in aiq_distribution/NetworkProxy/config/global.json so that Chrome browsers have following option added:

```
{
  "name": "experimentalDS3",
  "value": "--disk-cache-dir=cache|--disk-cache-size=1"
}
```

3. Start AIQ started using a systemd serviceDo not use /etc/rc.local.

The systemd service can set the maximum number of open files per process. This is needed in linux distros because the default of 1024 files per process is too low.

Note: Staring AIQ as a service in this manner is not compatible with the Lab Management automatic upgrade for linux nodes.

To configure the service create a text file named appvance.service in the /etc/systemd/system folder with the following contents. Edit the paths in file as needed to match your configuration.

```
[Unit]
Description=Start Appvance IQ services
After=syslog.target
[Service]
Environment=PATH=/opt/graalvm-ce-java11-21.3.3.1/bin:/sbin:/bin:/usr/sbin:/usr/bin
Environment=JAVA_HOME=/opt/graalvm-ce-java11-21.3.3.1
User=ec2-user
Type=simple
LimitNOFILE=10000000
LimitNPROC=20000
ExecStart=/bin/bash /opt/aiq/aiq_distribution/TestNode/nvme_setup.sh
WorkingDirectory=/opt/aiq/aiq_distribution/TestNode
TimeoutStartSec=0
[Install]
WantedBy=multi-user.target
```

After creating above file, execute the following commands in the shell:

```
systemctl daemon-reload
```

```
systemctl enable appvance
```

With the above configuration, an instance with size of m6id.2xlarge AIQ was able to execute a test designer test case using 40 chrome browsers concurrently. An instance with a size of m6id.4xlarge was able to execute 80 chrome browsers concurrently. Please note that specific values might vary depending of the application under test.

Last updated: May 26, 2023

If you have any feedback on this topic, please contact us at documentation@appvance.ai