

SPEC – Power and Performance

Design Document

SSJ Workload

SPECpower_ssj2008 V1.10

Standard Performance Evaluation Corporation

Table of Contents

1	Introduction	3
2	The SSJ Workload	3
2.1	Load Levels	3
2.1.1	Target Loads	3
2.1.2	Calibration	4
2.1.3	Active Idle	4
2.2	Transactions	5
2.3	Configurable Parameters	5
3	Director	6
3.1	Initialization	6
3.2	Runtime	7
3.3	Completion	7
4	Validator	7
5	Reporter	7
6	Raw File Format	8
6.1	SSJ Single JVM	9
6.1.1	Results summary properties	9
6.1.2	Miscellaneous per-run properties	9
6.1.3	Per-load-level properties	10
6.2	Director	11
6.2.1	Merging SSJ results	12
6.3	CCS	12
6.3.1	CCS Result Properties	13
6.3.1.1	Workload State	13
6.3.1.2	Per-interval properties	14
6.3.1.3	Per-interval power analyzer properties	14
6.3.1.4	Per-interval temperature sensor properties	15

1 Introduction

SPECpower_ssj2008 is the first industry-standard benchmark for measuring both the performance and the power consumption of servers. While some early attempts at describing the power consumption of servers simply measured the power consumed while running some standard performance benchmark, SPEC recognized that many servers include technologies to reduce the power consumption when the system is running at low utilizations. Since most systems spend much of their time running at less than full capacity, SPEC developed a methodology which advocated measuring performance and power consumption at a variety of system loads. SPECpower_ssj2008 is the first benchmark which implements this methodology.

The workload consists of a few main components:

- Server Side Java (SSJ): the workload itself
- Director: coordinates runs with multiple Java Virtual Machines (JVMs) and communicates with the Control and Collection System (CCS)
- Validator: verifies the configuration and results to identify runs which do not meet the SPECpower_ssj2008 run rules. (Note: Not all run rules are validated, and therefore some runs may be invalid even if they are not detected as such by the Validator.)
- Reporter: generates reports from the results

As described in the design overview, a full SPECpower_ssj2008 environment also includes CCS and PTDaemon. These utilities are not part of the SSJ workload itself, and are described elsewhere.

Perhaps the biggest new feature in SPECpower_ssj2008 v1.10 is support for running the workload across multiple operating system images at once, allowing for measurement of the power efficiency of multiple systems (such as blades) that share common infrastructure. The initial SSJ release was developed with this support in mind, so the main changes made to enable this support were in the reporter. Additional changes were made to make it easier to run in this type of environment.

2 The SSJ Workload

2.1 Load Levels

An SSJ run consists of two main phases: Calibration, and running at a series of Target Loads. The calibration phase is used to determine the maximum throughput that a system is capable of sustaining. Once this calibrated throughput is established, the system runs at a series of target loads. Each load runs at some percentage of the calibrated throughput. For compliant runs, the sequence of load levels decreases from 100% to 0% in increments of 10%. Measuring the points in decreasing order limits the change in load to 10% at each level, resulting in a more stable power measurement. Using increasing order would have resulted in a jump from 100% to 10% moving from the final calibration interval to the first target load, and another jump from 100% to Active Idle at the end of the run.

While the first phase of the workload is Calibration, the Target Load intervals will be described first.

2.1.1 Target Loads

At a high level, the benchmark models a server application with a large number of users. Requests from these users will come in at random intervals (modeled with a negative exponential distribution), and processed by a finite set of threads on the server. The exponential distribution may result in bursts of activity; during this time, requests may queue up while other requests are being processed. The system will continue processing transactions as long as there are requests in the queue.

The actual implementation of this model is a bit more complicated. The data set is composed of several mostly-independent “warehouses”. The driver schedules work independently for each warehouse, using Java’s `ScheduledExecutorService`. There is no physical “queue” of requests; instead, the driver maintains the “intended arrival time” of the current batch of transactions. Once the batch has been processed, the delay for the next batch of transactions is calculated, and this is added to the intended arrival time of the current batch to calculate the intended arrival time for the next batch. If this time is in the future, the next batch is scheduled to run at that time. If the intended arrival time is

in the past (i.e. if the server is currently running behind), the next batch will begin executing immediately. Since the intended arrival time of each batch is calculated based on the intended arrival time of the previous batch and the random delay, the end effect is the same as if the work actually queued up.

The delay between batches is calculated to achieve the desired throughput for each target load. Each target load runs at some percentage of the calibrated throughput.¹ (The mechanism of actually determining this calibrated throughput is described in the next section.) Transactions are scheduled in batches of 1000 transactions each on a per-warehouse basis. Before the pre-measurement period begins, the average delay between batches is calculated. This “mean delay” is the time between the beginning of one batch and the beginning of the next batch needed to sustain the desired throughput.

For example, if the calibrated throughput is 100,000 ssj_ops, the 40% target load would have a target throughput of 40,000 ssj_ops. With 8 warehouses, each warehouse needs to sustain a throughput of 5,000 ssj_ops. Since there are 1000 transactions per batch, each warehouse will execute an average of 5 batches per second. Thus, the mean delay between batches is 200 ms.

While the load is running, the actual delay between each batch is calculated using a negative exponential distribution using the mean calculated above, and capped at 10 seconds to prevent rare long delays from skewing the results. Thus, given a random value x in the range $[0.0, 1.0)$, the actual delay will be: $mean * -\ln(x)$, with a maximum of 10 seconds.

2.1.2 Calibration

At the beginning of the run, a series of calibration measurements are performed to find the maximum throughput of the server. During calibrations, transactions are scheduled and run in very nearly the same way as during the target load intervals. The only difference is that when a batch of transactions finishes, the next batch is scheduled to begin immediately, rather than having a delay between batches.

By default, 3 calibration intervals are used. The number of intervals is configurable, although the run rules specify a maximum of 10 calibration intervals for compliant runs. The calibrated throughput (used to calculate the target throughputs during the target load measurements) is calculated as the average of the last 2 calibration intervals. In this sense, all of the calibration intervals except the last two are really “warmup intervals”, since the throughput in these intervals is not used as part of the calibration calculation.

2.1.3 Active Idle

The intent of the Active Idle measurement is that the application is ready to accept requests, but that none are actually being received. To this end, the workload performs all of the same setup as is done for a normal interval (such as starting a scheduler thread for each warehouse). During the time of the interval, no transaction batches are scheduled, so no work will be performed.

However, note that CCS gathers status information for the workload once per second. This status information is sent even during idle, so the application isn't completely idle. This is representative of real environments where some level of monitoring is likely to be performed even when the system isn't otherwise busy; in practice, experiments have shown that this monitoring has very little impact on the utilization of the server.

¹ It should be noted that SSJ does NOT target a particular CPU utilization – it intentionally targets a percentage of the calibrated throughput. CPU utilization is not measured in a consistent way across different platforms, and can often be inaccurate in the presence of multiple threads per core or certain power management features. Targeting a percentage of the calibrated throughput ensures consistent behavior across platforms.

2.2 Transactions

SPECpower_ssj2008 executes 6 different transaction types, with the approximate frequency shown below:

- New Order (30.3%) – a new order is inserted into the system
- Payment (30.3%) – record a customer payment
- Order Status (3.0%) – request the status of an existing order
- Delivery (3.0%) – process orders for delivery
- Stock Level (3.0%) – find recently ordered items with low stock levels
- Customer Report (30.3%) – create a report of recent activity for a customer

Input data for each transaction is randomly generated. The transactions modify in-memory data structures representing Warehouses, Customers, Orders, Orderlines, etc. Nearly all transactions modify only the local warehouse – thus, there is minimal interaction between warehouse threads during the run.

2.3 Configurable Parameters

Only a small number of benchmark parameters are configurable for compliant runs. The most common of these properties are listed in the SPECpower_ssj.props file shipped with the benchmark kit. Modifications to these parameters will generally have little or no impact on the benchmark results; this ensures that results published by multiple vendors are comparable.

One property which may impact results is `input.calibration.interval_count`. This parameter specifies how many measurements are performed during calibration. The calibrated throughput is calculated as the average of the last two of these intervals. The default value of 3 intervals is sufficient for many systems, providing for about 5.5 minutes of warmup time (e.g. for the JIT compiler to compile key methods) before the real calibration begins. For some systems, a longer warmup time may be needed in order to reach steady state. This value can be increased up to a maximum of 10 for compliant runs.

For advanced users, the workload supports a variety of other configurable parameters which can be used for research and experimentation but not for compliant runs. These properties are described in the SPECpower_ssj_EXPERT.props file. Comments in that file describe each of the parameters. Some of the more interesting ones include:

- `input.load_level.count` – changes the number of load levels to measure. Rather than running 10 intervals at increments of 10% of the maximum throughput, you could run 20 intervals at increments of 5%, or 5 intervals at increments of 20%, etc.
- `input.load_level.target_max_throughput` – setting this value will override the calibrated value with a specific throughput target. The calibration intervals will still be measured (to allow the system to “warm up” before the measurements begin), but when you specify this value you can also use `input.calibration.interval_count` to set the number of calibration intervals to 0 or 1. Setting the maximum throughput manually could result in a more consistent set of results for a specific configuration. Or the throughput could be intentionally set too low or too high to see the system behavior at those loads.
- `input.load_level.throughput_sequence` and `input.load_level.percentage_sequence` – these properties can be used to manually specify a sequence of load levels, rather than the automatic behavior of starting at 100% and decreasing down to 10%. This could be used (especially in conjunction with a smaller value for `input.load_level.length_seconds`) to imitate a more dynamic environment where the system load is changing frequently in a more random manner. These parameters can also be used to insert idle intervals at any point during the sequence of measurements.
- `input.scheduler.single_queue` – as described in section 2.1.1, SSJ normally creates a `ScheduledExecutorService` for each warehouse. Each of these executors has its own thread “pool” (with only one thread) and its own virtual work queue. Many applications use a somewhat different model where a single thread pool is used to process all incoming requests. Setting `single_queue` to “true” will mimic this behavior, with a single thread pool (with the number of threads specified by `input.scheduler.number_threads`) and a common virtual work queue.

- `input.scheduler.batch_size` – by default, SSJ schedules 1000 fine-grained transactions in each batch. This value can be configured to make the workload more or less bursty in nature.
- `input.idle.*` – several properties can be used to change when and how active idle measurements are performed. For example, idle measurements can be taken immediately before or after calibration, instead of or in addition to the idle measurement at the end of the run.
- `input.ccs.enabled` – while CCS is an important part of the SPECpower_ssj benchmark suite, it is really only necessary when collecting power and temperature data. For performance-only measurements where power data will not be collected, CCS can be disabled; this can simplify the benchmark environment in non-compliant runs.
- `input.transaction_mix.*` – the transaction mix can be modified to change the characteristics of the workload somewhat. The various `transaction_mix` parameters can be viewed as the number of cards to put into a deck – the deck is then shuffled, and transactions are chosen from this deck. Note that changing the balance between transactions may result in growing data structures or other performance-related differences.

3 Director

The Director is responsible for keeping multiple SSJ instance in sync with each other, and for communicating with CCS as the run progresses.

3.1 Initialization

When Director starts up, it listens on a socket for incoming connections from each SSJ instance. One of the command-line parameters to Director tells how many unique hosts to expect connections from. As each SSJ JVM connects, it informs Director of the number of JVMs running on that host. Thus, Director can determine when the expected number of SSJ instances have connected. At this point, it waits for an additional 10 seconds for any extraneous connections, which most likely indicate a configuration error.

Once all connections with SSJ instances have been established, Director asks each SSJ JVM for the version of SSJ in use – if any versions do not match, the run is terminated.

Next, Director and the SSJ instances communicate to determine the offsets in the system time between the Director and each SSJ instance, as well as the approximate network latency. These values are used later by the Validator to determine whether all Recording intervals occurred while all workloads were operating at the same load level.

In the initial release of SSJ, each SSJ instance read run-time and system configuration properties from its own property file. With the ability to run SSJ on multiple hosts, this approach would make it easy for property files on different systems to be unintentionally out of sync, leading to invalid runs. To simplify configuration, the properties file can now be passed to Director, and Director will transfer the run-time properties to each SSJ instance. This is done after the clock offsets are determined. It is still possible to override these properties by providing specific property files to each SSJ instance; this could be useful for non-compliant runs where each SSJ instance may run a different sequence of load levels. The system configuration properties are also exchanged at this time; each SSJ instance has a “set id” which is used to identify the configuration properties file stored on the Director system. For example, if the set id is “sut” (the default), the configuration properties would come from the `SPECpower_ssj_config_sut.props` file on the Director system. Using different set ids can be useful for (non-compliant) heterogeneous runs, where the systems running SSJ are not all identical. Variable set ids can also be useful when a single Director system is used for a variety of different SSJ systems at different times – each type of SSJ system can use its own set id, and the configuration files on the Director system would not need to change from one run to the next.

After properties are exchanged, Director establishes a second socket connection to each SSJ instance to be used for status information. An additional connection is established with CCS. Throughout the run, Director obtains status information from each SSJ instance and aggregates it into an overall status which is provided to CCS.

3.2 Runtime

After initialization is complete, the run itself can begin. Director has two main jobs during the run. First, it communicates with the SSJ JVMs to align the start and completion of every load level. To accomplish this, each SSJ instance sends a message to the Director when it is ready to start/complete a load level. Director gathers all these responses. Once all SSJ workloads are ready, the Director broadcasts a start message telling all SSJ workloads to proceed.

The second task for the Director is to asynchronously aggregate status messages for CCS. This consists of several repeated steps. First, the director accepts and then broadcasts CCS queries to each SSJ workload. The responses to CCS consist of a set of properties. Upon gathering responses from all workloads, the director aggregates each property, in a prescribed way. Some property values must be identical for all SSJ workloads. Some property values must be summed across all SSJ instances. Some property values represent state information. If all workloads are at the same state the aggregation of that property is the value of any single property. If workload states are different, rules to aggregate to the correct state are applied. For example, if SSJ workloads are on different load levels at the same time, an error state is returned. Once all these properties are aggregated, they are sent to CCS.

3.3 Completion

At the end of the run, the Director gathers output files from each SSJ instance and generates a single combined raw file with the combined results from all JVMs. Director will invoke the Reporter to generate a performance report from these results, and forward the combined raw file and other data to the CCS system.

4 Validator

SSJ includes several validation tests designed to detect issues that will cause a run to be invalid. These tests are not intended to be comprehensive, but they will detect many common issues, and give the tester a high level of confidence that the result is valid. The specific tests that are performed are listed in section 2.5.2 of the Run Rules.

The validation tests are run automatically by each SSJ JVM at the end of the run, by the Director after it has combined the results from all SSJ JVMs, by CCS at the completion of the run, and by the Reporter when it is run. While many of these are redundant, running tests at all of these times ensures that the tester will see any validation errors even when looking at the results from only one component of the workload.

In addition to finding validity problems, the Validator can also generate warnings. These warnings are used to describe cases where it is possible or likely that a run is invalid or that configuration information is incomplete, but the Validator cannot prove it. If a run contains warnings, it is not marked invalid; however, the warnings indicate that the results may need further scrutiny to ensure their validity. Validator warnings will be shown in the report along with actual validation errors. They will be visible to SPEC during the review cycle; however, if a result is accepted by SPEC the warnings will not be shown in the report published on SPEC's website.

When a result file is validated, it is logically divided into two sets of values: ones which may be modified after the run (e.g. descriptions of the various components of the system) and the ones which may not be modified (e.g. the actual performance and power measurements). During validation, a checksum is computed for the un-modifiable properties. When the reporter runs, it verifies this checksum to ensure that none of these properties have been inadvertently modified.

5 Reporter

The reporter component is invoked automatically at the end of a SSJ run, and can also be invoked manually. The reporter is also used to generate reports for results submitted to SPEC. Regardless of how the reporter is invoked, its purpose is to read the raw results file from an SSJ run and generate one or more human-readable reports.

There are actually 3 variations of the raw file which the reporter must handle:

1. The raw file produced by a single SSJ JVM
2. The raw file produced by the Director for one or more SSJ JVMs
3. The raw file produced by CCS

The format of these raw files is described in the next section.

The reporter detects which type of raw file it is working with, and generates the appropriate type of report. The supported report types are:

1. The Power/Performance report (produced from CCS raw files)
2. A Power Details report (produced from CCS raw files)
3. An Aggregate Performance report (produced from Director raw files for multiple SSJ JVMs)
4. A Set Performance report (produced from Director raw files for multi-set runs)
5. A Host Performance report (produced from Director raw files for multi-system runs)
6. A JVM Instance Performance report (produced from individual SSJ JVM raw files, or Director raw files for a single SSJ JVM)

The reporter may generate multiple reports from a single raw file – for example, if a CCS raw file contains data from multiple SSJ JVMs running on multiple systems, the reporter will generate the main Power/Performance report, a Power Details report, an Aggregate Performance report, a Host Performance report for each host, and individual JVM Instance Performance reports for each JVM.

For any of these report types, the reporter begins by reading the raw file and summarizing the results. Next, the Validator is invoked to detect any problems with the run. In particular, the raw file checksum is verified to ensure that none of the “un-modifiable” properties have been changed from the original raw file contents. Any validation issues are described in the report.

While the reporter can generate both text and HTML output, the main report generation code is mostly format-independent; a generic in-memory representation of the report is built, and then a “renderer” is used to produce a specific report format. For practical reasons, there are a couple of exceptions to this approach – for example, the results header near the top of the report is laid out differently based on the format being used.

The reporter underwent many changes since the original release of SSJ in order to support multi-system results. Although only homogenous configurations are currently allowed by the SPECpower_ssj2008 run rules, SSJ and the reporter are capable of dealing with heterogeneous configurations as well, through the use of multiple “sets”. This capability may be useful for research and development purposes, even though it is not allowed for compliant results.

6 Raw File Format

As described above, there are actually multiple variations of raw files produced by SSJ, Director, and CCS. They differ primarily in their scope, and the prefix for property names.

Each of these raw files is generated as the run progresses (with a “.partial” suffix added to the file name). At the end of the run, the data is “redacted”. This process involves running validation, storing validation results and metric values in the raw file, organizing the keys in two sections (for properties that may be modified after the run and those that may not be edited), sorting each section in alphabetical order, and calculating a checksum on the un-modifiable properties to ensure that they are not edited. Redaction of the merged CCS raw file also prepares the file for possible submission to SPEC, by prepending a common prefix to all keys and normalizing the workload name so that all results use common key names.

The raw file is a simple text format consisting of name=value pairs. Thus, each line consists of a property name (conventionally consisting of multiple parts separated by periods), an equals sign, and a value which continues to the end of the line. Every key name is unique. Lines beginning with “#” are treated as comments. The order of lines in the file is unimportant; SSJ will generally write the file in a particular order to make it more human-readable, but changing the order of lines will not change the meaning of the file. In particular, reordering lines will not change the checksum which is used to detect inadvertent modification of non-descriptive properties.

6.1 SSJ Single JVM

Single JVM results are stored (by default) in the `ssj/results/SPECpower_ssj.<run_id>` directory, and have the name `ssj.<run_id>.<host_id>.<jvm_id>.raw`.

The only modifiable properties are the descriptive configuration properties (usually from `SPECpower_ssj_config.props` and `SPECpower_ssj_config_<setid>.props` from the Director system). These are copied directly from that file during run time. For example:

```
config.test.location=Warrenton, VA, USA
```

There are several categories of non-modifiable properties. This list is not exhaustive:

- A copy of the Java system properties for the SSJ JVM (for example, the `java.home` property). These are copied directly (i.e. no extra prefix).
- A copy of the SSJ input properties (typically from `SPECpower_ssj.props` from the Director system). These all begin with the "input." prefix, and are copied directly to the raw file.
- Results summary properties (below)
- Miscellaneous per-run properties (below)
- Per-load-level properties (below)
- Load-level summary properties (below)

6.1.1 Results summary properties

These properties are added to summarize the results of the run. For each of the 11 load levels there is a property:

- `submetric.ssjops.<interval id>` – the `ssj_ops` for this interval

The interval id is 0 for the first load level (100% for compliant runs) and increases to 10 for the last load level (active idle). Calibration intervals are not included in these intervals.

For compatibility with the previous version of SSJ there is also another property to indicate the `ssj_ops@100%` (which will match `submetric.ssjops.0`):

- `metric.ssjops` – the `ssj_ops@100%` for this JVM

6.1.2 Miscellaneous per-run properties

Several properties are added on a per-run basis:

- `run.cmdline` – the command-line parameters passed to SSJ (this does not include the JVM arguments)
- `run.cmdline.set_id` – the set id for this JVM (either passed on the command-line or the default value)
- `run.cmdline.host_id` – the host id for this JVM (either passed on the command-line or the default value, which is the hostname of the system)
- `run.cmdline.instance_id` – the instance id for this JVM (either passed on the command-line or the default value)
- `run.cmdline.jvm_instances_this_host` – the number of SSJ instances on this host (either passed on the command-line or the default value)
- `run.cmdline.property_file` – SSJ property file for this JVM (either passed on the command-line or the default value)
- `run.cpu.available_to_java` – the result of `Runtime.availableProcessors()`
- `run.hostname` – the hostname for this JVM
- `run.sequence_number` – the sequence number for this run
- `run.validity` – true if the results from this JVM pass all validity checks
- `run.validity.999_checkit` – true if the `ssj.jar` checksum was successfully validated
- `spec.benchmark.version` – the version of SSJ used in this run
- `spec.benchmark.version.date` – the date this version of SSJ was created
- `spec.test.date` – the date when the run occurred
- `ssj2008.spec.checksum` – a checksum of the un-modifiable properties in this raw file

- `result.validity.power_measured` – a boolean value indicating whether power was measured directly by SSJ (using `input.power_meter.enabled` – this will be false for compliant runs, since power is measured through CCS)
- `result.clock.offset` – the offset (in nanoseconds) between the clocks on the Director and this SSJ instance
- `result.clock.network_delay` – the maximum network delay (in nanoseconds) for communicating between Director and all SSJ instances
- `result.calibration.max_throughput` – the calibrated `ssj_ops` for this SSJ instance

6.1.3 Per-load-level properties

There is a group of properties added for each load-level. These properties are of the form “`result.test999.<property_name>`”, where “999” is a three digit number representing the interval. The first interval is 001, and the number is incremented for each successive interval. These intervals include all load levels, including calibration and idle. So a compliant run will normally include a total of 14 intervals, numbered 001 through 014.

Start and end timestamps are provided in both milliseconds and nanoseconds. These values are obtained using the `System.currentTimeMillis` and `System.nanoTime` APIs. The nanosecond values are the ones that are actually used for validation of interval lengths. The millisecond values are included in order to provide an absolute timestamp (nanosecond values are useful only for relative elapsed time calculations). On some systems (hardware, operating system, and JVM combinations), discrepancies may exist in the elapsed time calculated with these two different clocks; the Validator will generate a warning if it detects such a discrepancy.

- `result.test999.batch_count` – the number of transaction batches executed during this interval
- `result.test999.cust_report.count` – the number of Customer Report transactions executed during this interval
- `result.test999.delivery.count` – the number of Delivery transactions executed during this interval
- `result.test999.elapsed_nanoseconds` – the elapsed time (in nanoseconds) of the measurement interval (not including pre-measurement and post-measurement time)
- `result.test999.end_postmeasurement_time_nanoseconds` – the end timestamp of the post-measurement interval (in nanoseconds, useful only for comparison to other nanosecond timestamps in the same result)
- `result.test999.end_time_nanoseconds` – the end timestamp of the measurement interval (in nanoseconds, useful only for comparison to other nanosecond timestamps in the same result), not including pre-measurement and post-measurement time
- `result.test999.end_timestamp_milliseconds` – the end timestamp of the measurement interval (in milliseconds), not including pre-measurement and post-measurement time
- `result.test999.heapsize` – the total heap size (obtained with `Runtime.getRuntime().totalMemory()`) at the end of the post-measurement period
- `result.test999.heapused` – the used heap space (obtained with `Runtime.getRuntime().totalMemory()` – `Runtime.getRuntime().freeMemory()`) at the end of the post-measurement period
- `result.test999.interval_type` – the type of interval (calibration, load_level, or idle)
- `result.test999.label` – a descriptive label for the interval (e.g. “Calibration 1” or “80%”)
- `result.test999.max_batch_response_time_ns` – the maximum time (in nanoseconds) for processing a batch of transactions during the measurement interval
- `result.test999.max_warehouse_transactions` – the largest number of transactions executed during the measurement interval by any thread/warehouse
- `result.test999.min_batch_response_time_ns` – the minimum time (in nanoseconds) for processing a batch of transactions during the measurement interval
- `result.test999.min_warehouse_transactions` – the smallest number of transactions executed during the measurement interval by any thread/warehouse

- `result.test999.new_order.count` – the number of New Order transactions executed during the measurement interval
- `result.test999.order_status.count` – the number of Order Status transactions executed during the measurement interval
- `result.test999.payment.count` – the number of Payment transactions executed during the measurement interval
- `result.test999.score` – the `ssj_ops` for this interval
- `result.test999.short_label` – a very short description of this interval (e.g. “Cal.2” or “80%”)
- `result.test999.start_premeasurement_time_nanoseconds` – the start timestamp of the pre-measurement interval (in nanoseconds, useful only for comparison to other nanosecond timestamps in the same result)
- `result.test999.start_time_nanoseconds` – the start timestamp of the measurement interval (in nanoseconds, useful only for comparison to other nanosecond timestamps in the same result), not including pre-measurement and post-measurement time
- `result.test999.start_timestamp_milliseconds` – the start timestamp of the measurement interval (in milliseconds), not including pre-measurement and post-measurement time
- `result.test999.stock_level.count` – the number of Stock Level transactions executed during this interval
- `result.test999.target_throughput` – the target `ssj_ops` for this interval, or -1 for calibration intervals
- `result.test999.total_batch_response_time_ns` – the total time (in nanoseconds) spent processing batches of transactions. Note that since multiple threads will process transactions simultaneously, this value may be larger than the elapsed time.
- `result.test999.warehouses` – the number of warehouses (threads) used during this interval. This will be the same for all intervals.
- `result.test999.working_state` – the state of the workload which can be used to correlate with CCS intervals. This includes an interval id (using a different interval numbering scheme), interval type (“cal”, “lvl”, or “aidle”), and the suffix “_sum_”.

6.2 Director

Director will create its own raw file with some director-specific properties. At the end of the run, the raw files from all SSJ instances are gathered and merged with the Director raw file into a single combined raw file. The combined raw file results are stored (by default) in the `ssj/results/director/SPECpower_ssj.<run_id>` directory, and have the name `ssj.<run_id>.details.raw`.

The Director-specific properties begin with the prefix “director”. They include:

- A copy of the Java system properties for the SSJ JVM (for example, the `java.home` property is stored as `director.java.home`)
- `director.result.clock.network_delay` – the maximum network delay (in nanoseconds) for communicating between Director and all SSJ instances
- `director.result.clock.offset` – the offset (in nanoseconds) between the clocks on the Director and this system. For Director, this will always be 0.
- `director.result.test_method` – “Single Node”, “Multi Node”, or “Heterogeneous”, depending on the SUT configuration used.
- `director.config.<set_id>.units` – for each set (only 1 for compliant results), how many units are in that set. This property can be edited after the run.

The JVM-specific properties are merged using the algorithm described in the next subsection.

In addition, some overall run summary properties are added:

- `metric.aggregate_ssjops` – the total `ssj_ops@100%` across all JVMs
- `metric.ssjops_per_host` – the total `ssj_ops@100%` across all JVMs divided by the number of OS instances

- `metric.ssjops_per_jvm` – the total `ssj_ops@100%` across all JVMs divided by the number of JVM instances
- `submetric.ssjops.<interval id>` -- the `ssj_ops` for each non-calibration interval. The first interval (100%) has interval id 0, and the interval increments until the active idle interval with id 10.
- `run.validity` – true if the result passes all of the validity checks performed by the Validator
- `ssj2008.spec.checksum` – a checksum of the un-modifiable properties in this raw file

Finally, aggregate configuration fields are added:

- `aggregate.config.cpu.chips` – the total number of CPU chips across all hosts
- `aggregate.config.cpu.cores` – the total number of CPU cores across all hosts
- `aggregate.config.cpu.threads` – the total number of CPU threads across all hosts
- `aggregate.config.jvm.instances` – the total number of JVM instances across all hosts
- `aggregate.config.memory.gb` – the total amount of memory (in GB) across all hosts
- `aggregate.config.nodes` – the total number of nodes (physical hosts) in the run
- `aggregate.config.os.images` – the total number of operating system images across all hosts

The values of these fields are generated based on the configuration of the hosts in the result. If the host hardware configuration fields are edited after the run, the aggregate configuration fields must be edited to match. If the aggregate fields do not match, the Validator will report a warning.

6.2.1 Merging SSJ results

Director uses two strategies for merging results from multiple JVMs. For the actual results, each property from each JVM is copied directly to the combined raw file, with the prefix `<hostid>.<jvmid>`. Thus, “`result.test005.max_batch_response_time`” might become “`myhost.001.result.test005.max_batch_response_time`”.

For the configuration properties, Java properties, and input properties, the properties for multiple JVMs are merged when possible. For example, if all JVMs in a run have the same value for these properties (which is typically the case for nearly all properties in a compliant run), then they can be represented with a single global value.

If the property value is identical across all JVMs, it is added with the prefix “`global.`” If the property is not global but is common for all JVMs within a set, it is added with the prefix “`set.<setid>`.” If the property is not global or set-specific but is common for all JVMs in a host, it is added with the prefix “`host.<hostid>`.” If the property does not fall into one of the above categories (i.e. it is different for each individual JVM), it is added with the prefix “`jvm.<hostid>.<jvmid>`.”

Properties which are considered editable in their single-JVM forms are also editable in the merged forms. It is also allowable to turn modifiable JVM-specific properties into a single global property (similar modifications for host-specific or set-specific properties are also allowed).

6.3 CCS

At the end of the run, Director sends its combined raw file to CCS, where it is merged with properties from CCS properties to create a single raw file with all power and performance data from the run. This raw file is (by default) in the `../Results/ssj.<run_id>` directory and has the filename `ssj.<run_id>.raw` There are once again multiple sources of these properties:

- CCS runtime and descriptive configuration properties (from `ccs.props`) are copied directly with the prefix “`ssj2008`” (for example, “`ssj2008.ccs.config.hw.cpu`”)
 - As an exception, the CCS “workload” is always renamed to “`ssj`”. Thus, the “`ccs.wkld`” property value will always be “`ssj`”, and the “`ccs.wkld.<wkld_id>.*`” properties will always become “`ccs.wkld.ssj.*`”
- All of the properties from Director’s combined raw file are copied with the prefix “`ssj2008.wkld.ssj`”
- CCS result fields, defined in the next subsection

- `ssj2008.ccs.checksum.valid` – true if the CCS checksum is valid (i.e. the CCS code was not recompiled)
- Java properties for the CCS JVM are copied with the prefix “`ssj2008.ccs.`”
- `ssj2008.ccs.ptd.<device id>.Uncertainty` – a boolean value indicating whether uncertainty checking was in use for the given device
- `ssj2008.ccs.ptd.<device id>.*` – various other device properties obtained from PTD, not directly used by SSJ
- `ssj2008.ccs.ssj.version` – the version of SSJ which was used on the CCS system (to call the reporter)
- `ssj2008.ccs.version` – the version of CCS used
- `ssj2008.metric.performance_power_ratio` – the final SPECpower_ssj metric
- `ssj2008.ptd.<device id>.ptd.certified` – True if the device has been accepted by SPEC for use with SPECpower_ssj2008
- `ssj2008.ptd.<device id>.ptd.version` – the version of PTD used for communicating with this device
- `ssj2008.run.validity` – true if the result passes all validation checks performed by the benchmark
- `ssj2008.spec.checksum` – the checksum of the un-modifiable properties
- `ssj2008.submetric.performance_power_ratio.<interval id>` – the performance / power ratio for each non-calibration interval. The first interval (100%) has interval id 0, and the interval increments until the active idle interval with id 10.
- `ssj2008.submetric.power.<interval id>` -- the average power (in Watts) for each non-calibration interval.
- `ssj2008.submetric.ssjops.<interval id>` -- the `ssj_ops` for each non-calibration interval.
- `ssj2008.aggregate.config.*` – calculated in the same way as the aggregate config fields for the Director-level raw file

6.3.1 CCS Result Properties

CCS result properties have the form “`ssj2008.ccs.result.<interval_id>.<property name>`”. The interval id is an integer (with no leading zeroes) beginning with 1 and incrementing from there. The interval id does NOT match the SSJ interval id defined above. The interval id will be incremented each time CCS detects a change in workload state. Note that since some workload states are very brief, CCS may not detect the state change; thus, the interval ids may not match from one run to the next.

6.3.1.1 Workload State

The workload state normally consists of either one part (a run phase) or three parts: a load level id, the run phase, and the test state. The workload state begins and ends with an underscore (“_”) and the three parts are separated by underscores.

The run phase will be “cal” for calibration intervals, “lvl” for load levels, or “aidle” for active idle intervals. Each of these phases use the three-part form of the workload state. The remaining run phases are “error”, “init”, “inter”, “done”, or “settle” – these all use the one-part form of the workload state. Their usage is as follows:

- `error` – multiple JVMs have inconsistent states
- `init` – used during initialization
- `inter` – the time between the measurements
- `done` – the time after the run is complete
- `settle` – the settle time before active idle begins (not used for compliant runs)

The load level id (used in three-part workload states) is a three digit number (padded with leading zeroes) that normally starts at 1 and is incremented for each SSJ measurement interval (calibration, load levels, or active idle). The interval id is reset to 1 after calibration is complete. If a pre-calibration active idle period is used (not allowed for compliant runs), it will be interval 0.

The test state (used in three-part workload states) is one of:

- `error` – if the JVMs are in inconsistent states
- `inter` – for the time between measurement intervals (this state is used internally and will not appear in the raw file)
- `sum` – the summary at the end of a test interval
- `ru` – the pre-measurement (“ramp up”) period
- `rc` – the measurement (“recording”) period
- `rd` – the post-measurement (“ramp down”) period (this state is used internally and will not appear in the raw file)

The workload state can be used to link SSJ results with CCS results – the `result.test999.working_state` property from SSJ can be used to identify the SSJ interval corresponding to a particular CCS interval (`ssj2008.ccs.result.<interval id>.WK_STATE`) or vice versa.

6.3.1.2 Per-interval properties

There are many properties produced by CCS for each interval.

- `ssj2008.ccs.result.<interval id>.ACTION` – always “REQ_TX_SUMMARY”
- `ssj2008.ccs.result.<interval id>.CCS-ser` – provides a serial number for detailed CCS logging; in the raw file it will always have the value “### [wkld.<workload_id>]”
- `ssj2008.ccs.result.<interval id>.CCS_TIME` – the current timestamp
- `ssj2008.ccs.result.<interval id>.TotalWatt` – the total watts across all power analyzers
- `ssj2008.ccs.result.<interval id>.wkld.ssj.AVG_TXN` – the `ssj_ops` (transactions per second) during the interval
- `ssj2008.ccs.result.<interval id>.BATCH_COUNT` – the number of transaction batches executed during the interval by all SSJ JVMs
- `ssj2008.ccs.result.<interval id>.BATCH_RT` – the total batch response time (in milliseconds) during the interval by all SSJ JVMs
- `ssj2008.ccs.result.<interval id>.CCS_RT` – the average round-trip response time (in milliseconds) for communication between CCS and Director
- `ssj2008.ccs.result.<interval id>.TRANS` – always “-“ in the CCS raw file
- `ssj2008.ccs.result.<interval id>.WHSE` – always “-“ in the CCS raw file
- `ssj2008.ccs.result.<interval id>.WK_STATE` – the workload state in this interval

6.3.1.3 Per-interval power analyzer properties

For each power analyzer (with device id `<device_id>`), the following properties will be produced, each with the prefix “`ssj2008.ccs.result.<interval id>.ptd.<device_id>`”:

- `<prefix>.CCS_RT` – the average round-trip response time (in milliseconds) for communication between CCS and the PTDaemon instance for this device
- `<prefix>.NOTES` – the raw summary data returned by PTD during this interval (subject to change)
- Current data (`<prefix>.Amps.*`)
- Power Factor data (`<prefix>.PF.*`)
- Voltage data (`<prefix>.Volts.*`)
- Power data (`<prefix>.Watts.*`)
- Power measurement uncertainty data (`<prefix>.Uncertainty.*`)

For each of the measured data types (current, power factor, voltage, power, and uncertainty), the following properties are included:

- `<prefix>.avg` – the average value during this interval
- `<prefix>.badSample` – the number of error samples during this interval
- `<prefix>.goodSample` – the number of non-error samples during this interval

- `<prefix>.max` – the maximum value sampled during this interval
- `<prefix>.min` – the minimum value sampled during this interval
- `<prefix>.totalSample` – the total number of samples during this interval

6.3.1.4 Per-interval temperature sensor properties

For each temperature sensor (with device id `<device_id>`), the following properties will be produced, each with the prefix `"ssj2008.ccs.result.<interval id>.ptd.<device_id>"`:

- Temperature data (`<prefix>.Temperature.*`)
- Humidity data [when supported by the device] (`<prefix>.Humidity.*`)

The temperature and humidity data have the same sub-values as for power data (previous section).