


Measuring system performance


Why financial institutions need to carefully evaluate the cost versus the performance of payments technology




Executive summary

The payments industry today is a fast-paced, innovative environment where anything can happen. With increased transaction volumes and a proliferation of channels and payment instruments, the electronic payments space is more complex than ever. This exponential climb in the volume of cashless transactions puts added stress on the financial institutions (FIs) that process them and, if not managed properly, has the potential to be very costly in the future. FIs are under increasing pressure to make significant technology investments in order to meet the changing demands of the digital age. To future-proof their business and remain competitive in this dynamic industry, FIs need to assess the current systems they use to run their payments business and make the changes necessary to cope with increasing transaction volumes if they are to ensure current and future growth.

In this white paper, Compass Plus explains what an FI should look at when measuring system performance and how best to leverage this data to make an informed decision when reviewing software vendors.



Executive summary	2
Measurements of system performance	4
Read the small print	5
Metrics	6
Parameters	7
Cost	8
Conclusion	9




Measurements of system performance

To the consumer, the perfect payment interaction is the one they don't notice, however what they can't see behind the scenes can drastically affect their banking experience. System performance is therefore intrinsically linked to customer satisfaction. In this digital age, waiting minutes for a card payment to be accepted at a POS terminal or for an ATM to dispense cash is not acceptable when the customer knows that these transactions can be processed in seconds.

In a transaction processing environment these types of customer interactions are known as the end-to-end response time; which includes all of the processes that need to take place to fulfil an end-user action. Every customer interaction is an opportunity for a financial institution to demonstrate both their capabilities and differentiators, which is why transaction speed and system reliability are integral for business.

There are two main approaches to evaluate system performance: measurement and modelling. Measurement is the fundamental methodology utilised in the financial industry to benchmark transaction processing, as it is the most realistic means of evaluating performance. Modelling is the methodology used when initially building systems to theoretically test hypotheses. For the purpose of this white paper, all testing methodology referred to herein is implicitly regarding the measurement approach as the most true-to-life.



Read the small print

When evaluating the performance of payment software, confusion tends to arise when software vendors quote numerical data with no explanation behind what the data represents or the parameters within which these figures were collected. In the majority of cases this refers to the Transaction Per Second figure or TPS which, as a quotable statistic, is not as straightforward as it may seem and will be addressed in more detail later in this document.

Parameters for benchmark testing aren't standardised and it is often difficult or even borderline impossible to make comparisons between the results vendors are offering. Test environments are generally designed as low complexity; simulated under limited assumptions and on empirical evidence. This lack of clarity around the data and parameters used means that the verification of these statistics in real-time can be very problematic. The underlying solution to this problem is to understand the terminology presented and how these statistics relate to each individual FI.

There are three key areas for consideration by FIs when evaluating the performance of payment technology:

- **Metrics** – as we have already covered, TPS is the typical metric used to represent system performance. However, there is some debate around the definition of a transaction, therefore the type of transaction being tested and components therein need to be clearly defined by the vendor.
- **Parameters** – these directly depend on the test environment designed and the anticipated results. This is where things become complex. Variables can range from the type of hardware used (e.g. high end – high cost or lower end – lower cost), volume of equipment used (e.g. number of servers, database size), the performance capabilities of the hardware (e.g. the percentage capacity reached), to what is actually being tested (type of transaction), transaction volume (workload conditions, time of day, time of month, peak load, sustained peak load), and type of test (simulation, real-time scenario), etc. Simple environments often provide impressive figures that cannot be replicated outside of the initial simulation scenario.
- **Costs** – understanding how performance relates to cost is imperative so as to measure this against the FI's specifications. Vendors should offer FIs options to increase transaction volumes incrementally, scalable to their business growth. To buy software that can offer substantially higher performance than required at the offset is not necessarily a financially sound decision.

FIs need to look at the bigger picture and read the small print when evaluating technology critical to the future success of their business.

As discussed above, when it comes to evaluating a vendor's payment software it is important to define and explain the performance metrics they provide to demonstrate their system performance. Without transparency in performance metrics, FIs will find it near-impossible to begin evaluating the real value of the payment technology. If business decisions are made without having a clear understanding of what the data represents, FIs may find themselves disappointed with the end result.

The key metric offered to illustrate system performance is the TPS. Whilst this definition may seem straightforward, the interpretation of a transaction can vary substantially from vendor to vendor. There are no agreed standard classifications or industry standardised definitions of performance metrics and, as such, it can be difficult to understand what the results of benchmark tests are demonstrating in actuality.

Transactions are essentially collections of actions that can be broadly classified as either business transactions or database transactions, and there is a fundamental difference between the two. A business transaction is a customer interaction in the 'real world'; withdrawing money from an ATM or making a purchase at a point of sale. These can also be referred to as nested transactions whereby some of the steps involved in making the transaction are other transactions, referred to as sub-transactions or, for the purpose of this paper, database transactions. A single business transaction could be the equivalent of 15 database transactions. Database transactions are therefore essentially the number of processes that happen behind the scenes to enable the overlying business transaction.

For example, transferring funds between two bank accounts is one business transaction, however, behind the scenes, the source account is debited, the target account is credited, this action is recorded, checks are made on the source account funds availability, etc. These sub-transactions account for the 'behind the scenes' database transaction metrics.

There is no wrong or right in how a vendor decides to represent a transaction when referencing system performance. The importance lies in the FI understanding the type of transaction being measured and subsequently what this number illustrates. Using TPS as a comparative measurement of system performance when one vendor is demonstrating database transactions and the other is showcasing business transactions, is essentially like comparing grains of sugar to a sugar cube and will not result in an accurate representation of the two systems against each other.

Vendors should be happy to fully disclose the details of the test environment that has led to the published results; providing prospective customers (FIs) performance information beyond the summary of the performance metrics. These should include statistics such as average authorisation time, percentage stability through peak load, time peak load was sustained and the percentage capacity of the server. Many vendors abstain from sharing details around how long peak load was sustained during stress testing. With intensive transaction periods becoming more common, there is an increased requirement for an FI's payment infrastructure to remain stable at higher capacities for longer.

Similarly figures that are presented based on the assumptive multiplication method, with terminology such as 'the equivalent of' or 'this equates to' can be misleading. Vendors multiplying TPS results to demonstrate the system's capabilities over an hour aren't realistically representing performance, especially if they have no solid evidence to demonstrate that peak performance can be maintained for that time period.

Performance testing is essentially used to determine how a system performs in terms of responsiveness and stability under a particular workload, whilst investigating, validating and verifying other attributes of the system such as scalability and reliability. It is therefore clear that measuring and promoting system performance using only TPS as an indicator is essentially redundant in understanding the mechanisms involved in generating these figures.

Parameters

The parameters for benchmarking system performance aren't uniform; there are multiple variables at play. These variables make up the test environment and it is essential to understand the components of this environment in order to understand what is being evaluated. Ideally, the parameters of any system testing should make up a simulation that is as close to reality as possible; to provide performance indicators that would be achievable in real-time. However, this is not always the case; often test scenarios are simplistic and subsequently unrealistic.

Ultimately, the test environment's primary objective should be to help the vendor better understand the limits of their system. By stress testing the system, the vendor can uncover and solve problems before go-live. The test environment should help create better technology rather than provide a 'version' of performance that is favourable to the vendor.

It is not always possible to stage a test environment that accurately represents the variables that occur in real-life. However, test results that demonstrate different volumes of different types of data (for example issuing and acquiring data) can replicate how the system behaves in various situations. Parameters that should be clearly outlined in the test scenario include the hardware used, the volume of transaction data used, types of data sets (for example, overdue/non-overdue accounts, active/blocked cards, etc.), the number of customers, cards and accounts that the test processing system database holds respectively, the number of terminals, types of terminals, protocols used, the day (for example, statement day, end-of-day, etc.), amongst others.

The combination of these types of parameters not only offer a close to real-time test simulation, but also offer the vendor data with which to better forecast server hardware requirements, and therefore provide savings for their current and future customers' transaction processing business.

Employing the methodology of stress testing the system by simulating real-life production-use as opposed to straightforward transaction switching is, in many respects, unique and is only performed by a handful of vendors. However, this is the only true test that expressly shows how the system will fare in a real-life situation.

Transaction processing systems should guarantee the four ACID properties:

Atomicity: Failed transactions do not produce side-effects: either a transaction completes its action on the database, or it has no effect.

Consistency: A transaction moves the database from one consistent state to another.

Isolation: A transaction depends on other transactions only through the changes they make to the database.

Durability: After a transaction completes, the changes it makes to the system state persist indefinitely.

It is very difficult to compare the price and performance of the various systems available in the marketplace today. Highly effective IT systems are essential to support growth, the future diversification of products and new channels, and the optimisation of an FI's infrastructure; and this can come at a cost.

Understanding the wider context of the test environment is integral to deciphering potential overheads. Metrics can look compelling at face value but the most impressive level of performance will always have specific hardware or maintenance requirements. For a payments system to achieve a certain TPS metric it will require a definitive number of servers or System Admins, for example, which may represent significant hidden costs.

The most effective measure of system performance for an FI is the Cost Per Transaction or CPT. This takes into account the hardware, software and operational costs; essentially the total system cost including maintenance, normalised by the performance metric. Although TPS is an important measure of system performance that should not necessarily be overlooked, it is the CPT that represents the system efficiency in direct relation to outgoing cost.

To ensure the best possible CPT, a vendor must demonstrate how its payment software will behave in a variety of real-life situations in order to propose an optimal system and hardware configuration based on the customer's requirements.

Bank systems are complex, made even more so from bolt-ons and plug-ins often using legacy systems that weren't built for the sheer volumes and transaction loads that are typical of today. Though overhauling the system an FI runs on might seem extreme and may not be necessary on a case-by-case basis, it is something traditional brick and mortar institutions will certainly need to consider in the next 10 years. The advantages of modernising your payments infrastructure include increased flexibility and scalability in an environment that is altogether easier and cheaper to run and manage in the long term.

Considerations when assessing performance versus cost for evaluating current systems and/or weighing up system replacement:

- Establishing the hidden costs: how many servers are running? How much do they cost? How many System Administrators are required for maintenance?
- Meeting current and future requirements: is the system scalable? How much will you need to pay for additional servers and support staff as your company grows? Are costs incremental over time?
- Keeping cost relative: what systems do you need to run your day-to-day business? Will the value of a system upgrade offset the hardware/software/maintenance costs? Does your business need to run on high-end large scale systems or would something smaller be more efficient?

Conclusion

Evaluating payment systems is difficult and complex for FIs. Whilst analyst firms do their best to compare software vendors based on the limited information made available to them, it cannot be guaranteed that the same metrics on which they have both weighted and used to base their analysis on are completely in line with the requirements of all FIs.

There is a lack of information, guidelines and advice to help FIs understand what supplier performance metrics really mean to their organisation. As a result, the majority of FIs simply don't evaluate in terms of performance versus cost and have a view of performance requirements that are neither realistic nor holistic. They risk paying excess for a payment system that does not meet their needs right now, and spiralling incremental costs to meet performance requirements in the future.

To avoid these pitfalls in evaluating new investments, FIs must read the small print and ask the right questions when making technology choices. The best approach is to clearly define what is being measured (metrics), what it is measured against (parameters) and what the cost will be, whilst not viewing technology in isolation.