

Business application monitoring

How to ensure the performance of your critical applications



Contents

CHAPTER 1
Business applications and their
technical environments

CHAPTER 2
Business applications, mission-
critical tools

CHAPTER 3
How can you make sure your
business apps perform well?

3 Introduction

4-9 Business applications and their usage
contexts

10-12 Ways of bringing business applications
to users

13-16 Consequences of poor application
performance

17-18 Quantifying potential losses

19-22 Causes of poor performance

22-24 Pre-production testing: regression
and load

25-31 Monitoring business applications
in production

32 Conclusion

Introduction

22

minutes¹. **That's how much time people waste each day because their computer is slow or the critical applications their company uses aren't working the way they should.** Just think of the costs that result from a loss of productivity! 22 minutes each day amounts to 91 hours per year! Imagine what people could do with that wasted time.

"Chat with their coworkers over coffee? Take care of the plant that's sitting in a corner of their office? Tidy up their desk or sort through e-mails?"

The performance of business applications, rightly called 'mission-critical', is a topic that speaks directly to a number of parties: **CTO, IT, CIO, business application providers, application managers, corporate business lines,** and more.

With the trend toward digitization of all processes, **good performance and reliable operation of business applications** is more vital than ever. They have an impact on the productivity and efficiency of the entire enterprise.

A wide variety of factors can affect the performance of these tools that are vital the daily business of the organization:

- version upgrades,
- changes in data centers,
- migration to the cloud,
- diverse updates,
- production incidents.

This being so, it is important **to be able to prevent performance deteriorations and failures by detecting them early on**, if possible before they affect users, and take steps to resolve them upstream.

This white paper tours the landscape of business applications and looks at their criticality for enterprises. It reviews some of the challenges that IT managers and CTOs are facing; these same challenges are a central concern for their providers of hosting and SaaS solutions. Next, we take a look at the possible causes of degraded performance and how to remedy them.

We hope you will benefit from our white paper!

¹ Source: Robert Half Technology

Business applications and their technical environments

CHAPTER 1



What are business applications? In what context are they used? What type of installation is best?

Business applications, indispensable for the operation of the enterprise as a whole, are involved at every level of the company.

1. Business applications and their usage contexts

Whether they housed within the company (**on premise**) or delivered as a service (**SaaS**), business applications are present at **every level of the value chain**.

Their primary objective is to implement and automate the enterprise's management processes to enable employees to perform their tasks as efficiently as possible.

Types of application

Business applications, as their appellation indicates, are business-specific; they are also sector-specific (banking, energy, insurance, automobile, etc.), and process-specific; they are, in most cases, customized to match the particular characteristics of each company.

Among the most popular major types of enterprise software are, for example, **ERP** (Enterprise Resource Planning), **CRM** (Customer Relationship Management) and **SCM** (Supply Chain Management).

Among the software the most specific to production, **CAD** (Computer Assisted Design) and **PLM** (Product Lifecycle Management) are some of the most widely used.

Naturally, **messaging/e-mail and online collaboration tools**, especially Microsoft's Office 365 package, are used by a growing number of enterprises.

ERP (Enterprise Resource Planning)

ERP is the production data management tool for all services and departments, from the supply chain all the way through sales:



Supply Chain

With SCM (Supply Chain Management) software, the enterprise can optimize how it handles orders, procurement, and inventory, as well as customer service.

Production / Manufacturing / Inventory / Fulfillment

With a Manufacturing Execution System (MES), the company has visibility over production processes and can oversee the factory workshop (shop floor).

HR (Human Resources)

Human resources software automates and simplifies HR management. Processes like payroll are optimized, and employees can schedule holidays and days off, handle expense reports, or request training.

Finance

Financial software digitizes accounting, with bookkeeping of payments and invoicing, and handles assets as well as processing and recording of all financial transactions.



The different types of software listed above are often interconnected and operate with each other in real time.

CRM (Customer Relationship Management)

CRM **records, processes and stores all information about customers and prospects** (the contact's particulars, specific needs, etc.).



CRM also interacts with other software, such as e-mail applications and marketing automation tools (management and automation of marketing tasks in line with the user's needs, including personalized e-mailings).

Any of the enterprise's units which has access to the CRM can view the history of conversations with the customer (retail, customer support, etc.) and the actions that were taken. CRM makes for improved follow-up over time and better sales performances.

By 2021, enterprises will devote the greatest share of their software spending to CRM technology, a Gartner study found. And by 2025, the CRM market should rise as high as USD 81.9 billion, which amounts to a yearly increase of 12% on the basis of a market worth USD 26.3 billion in 2015.

ITSM (IT Service Management)

ITSM tools such as ServiceNow enable IT staff to efficiently handle large-scale processing of internal and external IT requests (technical support, helpdesk, etc.).

E-Commerce

The back-end of online commerce is managed by applications like Magento, SAP Hybris, Episerver, to name just a few.

Multi-channel services like "click to collect" (go to a store to pick up an order placed online) rely on solutions that may be external, such as SoCloz.

Users and usage contexts

Company employees are the primary users of business applications, and they can be just as demanding as consumers visiting an e-retail site. People in the office expect robust, adaptable work tools. Providers or subcontractors may also have access to those business applications.

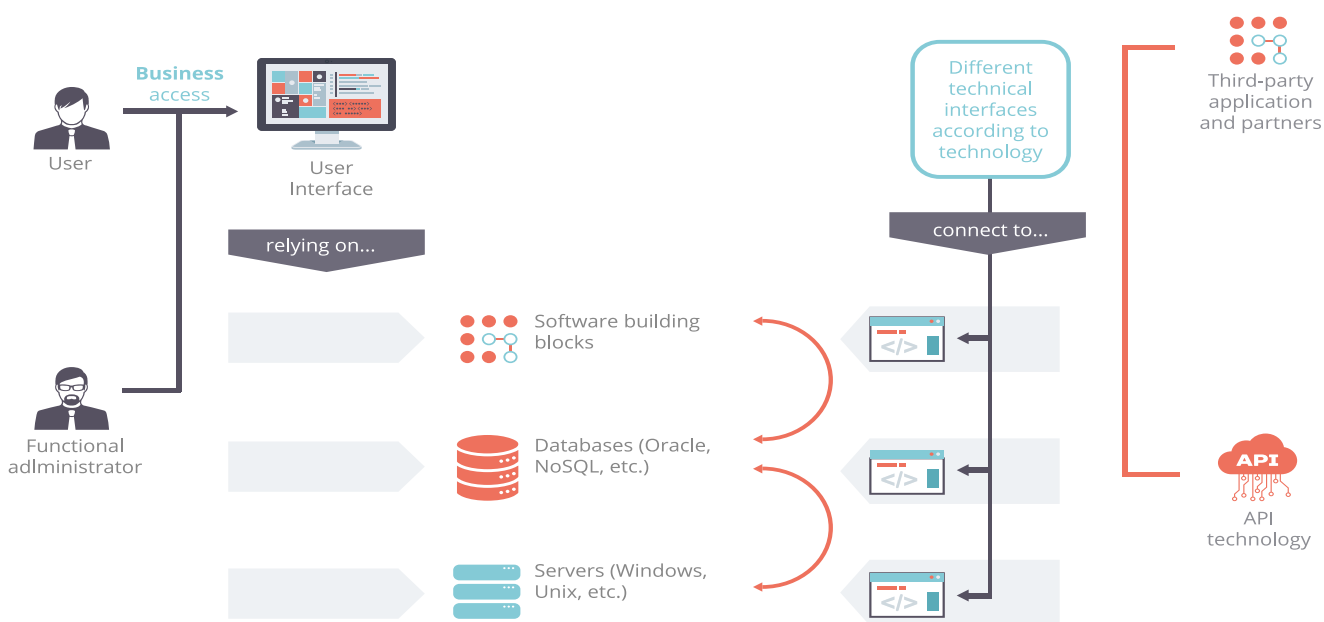
Furthermore, 'in-house' users need access to such applications whether they happen to be at headquarters or another corporate location, when they are working from home or on the road, at production sites, factories, on construction sites, or elsewhere. A car dealership or an insurance company, for instance, must ensure that its workforce around the country can log on with their tablets to the application they need to interact with customers.

It is a fact that a growing number of users connect to business applications **from mobile devices – tablets or smartphones. 48% of marketing specialists now connect to their CRM platform on their mobile phone.**

Of these same specialists, 87% consider that mobile applications help improve sales figures and increase workforce productivity.

Complex architectures

The end-to-end performance of any enterprise application involves a great many components – internal and external, network and application – any of which may be at the origin of service disruptions that can be tricky to detect.



Because managing the different components, layers, and parts of the architecture can be a complicated affair, the corporate trend is to host applications in the cloud to avoid becoming bogged down with the physical components (servers, data center, etc.).

Thick client, thin client, virtualization, cloud... and the end use

Business applications can be delivered to the user in a number of ways:

- **Web application:** this is an increasingly popular choice
- **Thick client,** Mainframe...
- **Thin client:** virtual desktop or virtualized application (Citrix, AWS AppStream, etc.)

While migrating to the cloud is one way enterprises have adopted to rationalize their IT, virtualizing applications (virtual desktop, etc.) is another, which may (or may not) be combined with a move to the cloud. Business application virtualization is very widespread because it streamlines production rollout and updating on fleets of devices (PCs).

However, **virtualizing applications means dealing with an extra software layer**, and this can be a real performance headache for users and operators. This is why it is important to **be able to accurately measure how the apps are performing** on the end-user side to make sure user productivity is not adversely affected.

Now let's take a closer look at some of the challenges IT executives face when deciding how to provide business applications to their enterprise in keeping with budget constraints, available human resources/skills, and corporate policy.

2. Ways of bringing business applications to users

There are a variety of configurations, each with the particularities and complexity of their given circumstances. This white paper takes a look at three such ways (**on premise, SaaS, and external IT providers**).

Installation on premise

Despite the attraction of the cloud, sometimes enterprises **need to deliver applications from their corporate data center**, particularly in business sectors like banking, defense, and others which require a high level of data security, sometimes without access to the internet. Furthermore, some applications need to be deployed on powerful workstations with a GPU for 3D processing (CAD, for instance) or machines with massive computation capabilities.

It may also be the case that no off-the-shelf software package is suited to the enterprise's needs, so the choice is made for on-premise software that is customized for a specific business activity or for the specificities of the enterprise's architecture (interconnection with legacy applications, etc.).



Advantages

- The possibility of adapting the application or adding additional application layers.
- Data is stored locally (confidentiality to shield against incompatible data protection laws of other countries, for instance).

Limitations

- Managing applications internally takes a lot of manpower
- Deployment and updates are labor- and time-consuming



SaaS (Software as a Service)

Year after year, the cloud is becoming more appealing to more organizations.

A couple of years ago, Sid Nag, research vice president at Gartner, said the trend was only just beginning. "Through 2022, Gartner projects the market size and growth of the cloud services industry at nearly three times the growth of overall IT services."

Gartner forecasted that the total cloud market would reach **USD 331.2 billion in 2022**.

Table 1. Worldwide Public Cloud Service Revenue Forecast (Billions of U.S. Dollars)

	2018	2019	2020	2021	2022
Cloud Business Process Services (BPaaS)	45.8	49.3	53.1	57.0	61.1
Cloud Application Infrastructure Services (PaaS)	15.6	19.0	23.0	27.5	31.8
Cloud Application Services (SaaS)	80.0	94.8	110.5	126.7	143.7
Cloud Management and Security Services	10.5	12.2	14.1	16.0	17.9
Cloud System Infrastructure Services (IaaS)	30.5	38.9	49.1	61.9	76.6
Total Market	182.4	214.3	249.8	289.1	331.2

BPaaS = business process as a service; IaaS = infrastructure as a service; PaaS = platform as a service; SaaS = software as a service

Source : Gartner



Advantages

- no time-consuming installation or updates need to be conducted by internal staff
- greater pricing flexibility, as in most cases the software is available by subscription with the possibility of cancellation subject to conditions, as well as the possibility of sliding-scale pricing by number of users
- SaaS is a business operating expense (OPEX), unlike CAPEX investments in software installed on the company's own servers and amortized over several years

Limitations

- the client enterprise is dependent on the service provider
- the SaaS vendor and its service providers are responsible for data management, security, and storage, often a sensitive point for enterprises
- changing providers is cumbersome because data transfer can be a hassle
- loss of control over the application
- interoperability problems may arise between the SaaS application and the company's internal IT systems



External IT providers (outsourcing)

Without going as far as SaaS, enterprises may opt to **delegate the management of its IT systems, equipment** (fleet of computers), **and business applications to an external provider**. Up to now, this has been the most usual approach that companies have taken.



Advantages

- Significant manpower savings (rationalization)
- The enterprise can focus on its own core business

Limitations

- The client enterprise entrusts its IT management to a third party, thereby losing control over business applications and their data
- The client enterprise is dependent on the outsourcing service provider, implying a relationship of total trust.
- Complex change management



A survey conducted by Markess International found that 37% of participating IT decision-makers feel that the ROI (return on investment) achieved or anticipated by moving their IT infrastructures to the cloud is higher than the ROI for conventional outsourcing, including IT service companies. Furthermore, the survey found that 50% of those decision-makers emphasized the gains in flexibility, time, and agility offered by cloud implementations over conventional outsourcing solutions.

In all cases, the loss of control and the lack of specified remediation associated with IT outsourcing (SaaS, service company, integrator) are not merely incidental risks. This is why contracts must include the following specifications, among others:

- **SLA (Service Level Agreement)**, stipulating the quality of service expected
- **MTTI (Mean Time To Intervention)**, the time from when the incident is registered to the first intervention
- **MTTR (Mean Time To Resolution)**, the time between when the incident is registered and its full resolution

To underscore the importance of guaranteeing service levels, we have only to consider the criticality of business applications and the financial impact of downtime. Let's look inside what the business app means to the enterprise...

Business applications, mission-critical tools

CHAPTER 2

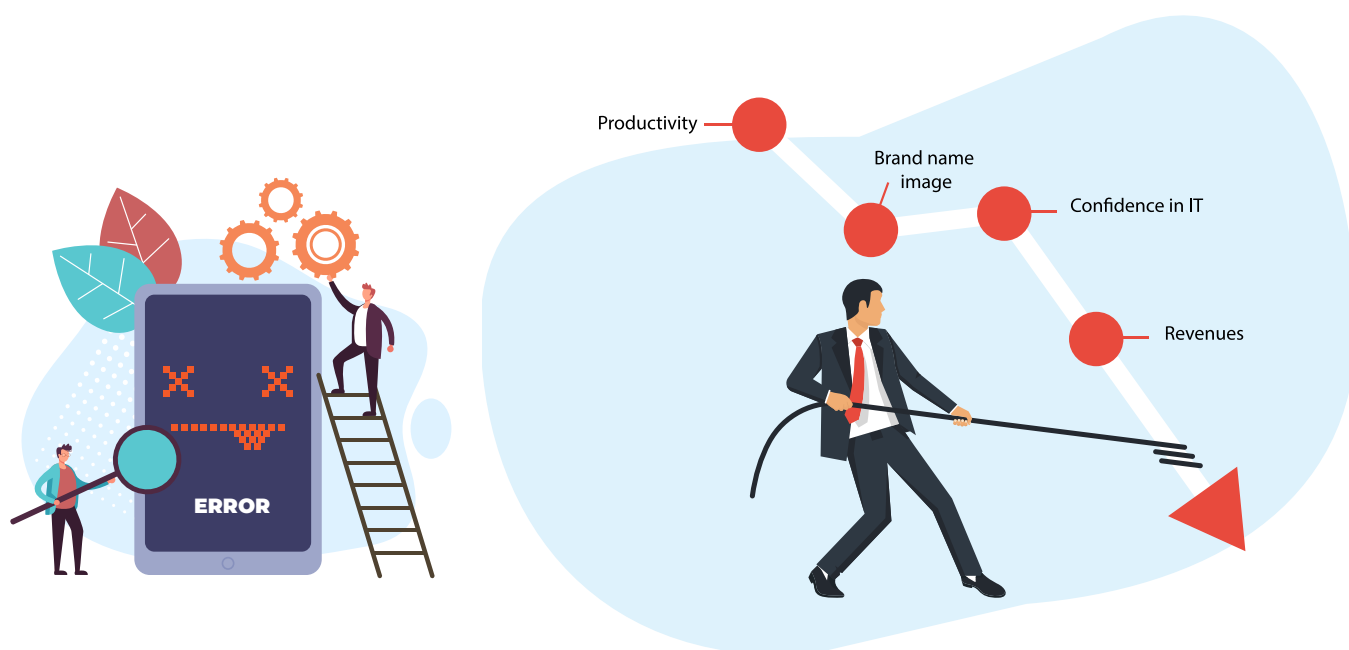


What consequences does degraded performance have? How can losses due to business application incidents be calculated?

These applications, used throughout the organization across all departments, are a determining factor of the enterprise's business continuity.

1. Consequences of poor application performance

Unavailability or poor performance from an application, whether the problem is occasional or recurrent, can have negative effects on the enterprise. As soon as critical applications become sluggish or fail to respond, productivity is affected, brand image is tarnished, the work atmosphere deteriorates, and business suffers.



The impact on productivity

Applications and software suites, thin client or thick client, are used on a daily basis by the company's workforce. People obviously cannot do their work without their work tools. If an application is unavailable or does not react properly to user requests,

serious difficulties may arise. The financial department cannot tolerate any delay in processing certain affairs; nor can human resources or sales teams. Downtime can even go as far as interrupting production lines.

22 minutes, That's how much time people waste each day because their computer is slow or the critical applications their company uses aren't working the way they should, according to a study conducted by Robert Half Technology. 22 minutes each day amounts to **91 hours per year** on the basis of 8 working hours per day for 50 weeks!

Just think of what people could do with that wasted time.

Chat with their coworkers over coffee? Tidy up their desk or sort through e-mails? Take care of the plant that's sitting in a corner of their office? Brainstorm?

"Technology is a double-edged sword in most organizations -- it can be an enormous time-saver, but it can also be a drain if things aren't working well or people don't know how to use the tools provided to them," says John Reed, executive director of Robert Half Technology. "A proactive and highly proficient help desk can be a huge productivity asset to companies and their employees."

John Reed, executive director of Robert Half Technology.

Loss of trust in corporate IT or external application provider

When an application malfunctions or is unavailable, users complain to their company's IT department. And yet the technical staff might find that everything appears to be running smoothly on the infrastructure.

The legitimacy of the IT department may take a beating as internal audiences lose confidence in technical staff.

This paves the way to general discontent towards the IT department which is held responsible for ensuring that business applications operate smoothly.

If application services are outsourced, the IT department turns to its application provider for **the reasons behind poor performance**. The causes can be located at any link in the delivery chain.

Sometimes the difficulty of pinpointing problems gives rise to misunderstandings, each side believing the other is not fulfilling its commitments.

If no objective, pertinent data is available to both parties, there can be no basis for communication to solve problems.

Deteriorated brand image

It goes without saying that downtime or unavailability of business applications can have very negative consequences on the image customers have of the enterprise, either directly or indirectly.

Any application performance problem which affects the customer relationship, from the application which records orders and every step up to invoicing, to enterprise telephony and voice applications, is

obviously incapacitating. If the CRM software freezes while an employee is trying to respond to a customer's request, that's a clear sign of inefficiency.

The risk is always that **customers will turn to a competitor that is perceived as being more efficient**, with better performing technology.

Financial impact

Business applications have a direct effect on productivity, and therefore on the enterprise's profitability:

- Downtime of mission-critical applications is a **waste of time** for employees.
- Service disruptions on customer-facing retail applications mean a **loss of revenues**.
- In the case of an enterprise which has signed an SLA with customers, it must deliver the specified service levels; otherwise it must compensate its customers by **paying penalties**.
- Incidents translate into **additional IT spending** on problem resolution.

Below we take a detailed look at the potential costs of degraded performance.

2. Quantifying potential losses

Dr. Bill Curtis, founding executive director of the Consortium for IT Software Quality (CISQ), estimated the potential cost of degraded application performance. He defines this cost in terms of both a **reduction of productivity** and a **loss of revenues**. This principle is illustrated in the formulas below:

$$\begin{aligned} \text{Loss} = & ((\text{avg. revenue per minute}) \times (\text{number of minutes unavailable})) \\ & + \text{costs to reactivate business} + ((\text{Additional customer service minutes}) \\ & \times (\text{€ per minute})) \\ & + \text{future revenue lost from defecting customers} \\ & + \text{liquidated damages if applicable} \\ & + \text{other related costs}) \end{aligned}$$

$$\begin{aligned} \text{Loss} = & [(1 - (\text{reduced output under degraded performance} \div \\ & \text{avg. output at normal performance})) \\ & \times \text{number of workers affected} \\ & \times \text{avg. hourly cost} \times \text{hours of degraded performance}] + \\ & [(1 - (\text{reduced output or revenue under degraded performance} \div \\ & \text{avg. output or revenue at normal performance})) \\ & \times \text{performance} \\ & \times \text{avg. revenue per hour}] + \\ & \text{costs of delayed work} \end{aligned}$$



Example

Take an application that generates transactions worth 200 euros/minute (12,000 euros/hour). Suppose that the application breaks down for 3 hours. What is the impact on the enterprise? How much revenue is lost? What is the total cost of the failure?

After noticing that the application is not responding, three of the enterprise's employees work to locate the problem affecting the application, fix it, and test the application after resolving the incident. This takes them 3 hours of work each, at 20 euros per hour. The outage also results in an increase in calls to customer service, representing 150 additional minutes of calls at 2.50 euros per minute.

$$\begin{aligned} \text{Loss} = & [(200\text{€ per minute}) \times (180 \text{ minutes})] \\ & + [(3 \text{ hours} \times 3 \text{ employees}) \times 20 \text{ € per hour}] \\ & + (150 \text{ minutes of additional calls} \times 2.50 \text{ € per minute}) \\ & = 36,000 + 180 + 375 \\ & = 36,555 \text{ €} \end{aligned}$$

Supposing that the enterprise experiences three such outages in a year, the annual cost reaches 109,665 €!

Next consider that 300 in-house users need this application for their everyday work. These users are paid 18 € an hour. Also consider that the application's performance drops by 5% because of database access issues. This loss of productivity is calculated as a loss both of revenues and the productive value of the employee's salary.

The cost is estimated at:

$$\begin{aligned} \text{Loss} &= [(0.05 \text{ reduced output}) \times (300 \text{ affected} \\ &\text{employees}) \times (18 \text{ € per hour per employee}) \times (1600 \\ &\text{hours per year})] + [(0.05 \text{ reduced output}) \times (12,000 \text{ €} \\ &\text{revenue per hour}) \times (1600 \text{ hours per year})] \\ &= 432,000 + 960,000 \\ &= \mathbf{1,392,000 \text{ € per year, or } 348,000 \text{ € per quarter}} \end{aligned}$$

The annual cost of the application's quality problems is therefore extremely consequential for a business. In these calculations, we did not count potential penalties for delayed work within the company. So the cost of application outages can be even higher.

It is therefore crucial for any company to **make sure its business applications are performing well** and **monitor the application over time to address the risk of failures and incidents.**

How can you make sure your business apps perform well?



Given the complexity of application delivery chains, how can you make sure everything is going well? How can you know whether user complaints are justified? How can you check whether each party is complying with performance objectives?

If you don't collect any metrics, you're flying blind. If you collect and focus on too many, they may be obstructing your field of view.

Scott M. Graffius,

Agile Scrum: Your Quick Start Guide with Step-by-Step Instructions

Everything is measurable; the problem is knowing what to measure. As Scott M. Graffius cautions, “If you don’t collect any metrics, you’re flying blind. If you collect and focus on too many, they may be obstructing your field of view.” You must know exactly what to measure if you want to make headway, and the choice is vast: networks, servers, client PC, user experience, application performance, and so much more.

A recent study conducted by Gartner found that **61% of respondents identified APM** (application performance management) as being **important or crucial**. Of those, **46% stated that monitoring end-user experience is the most critical dimension of APM.**

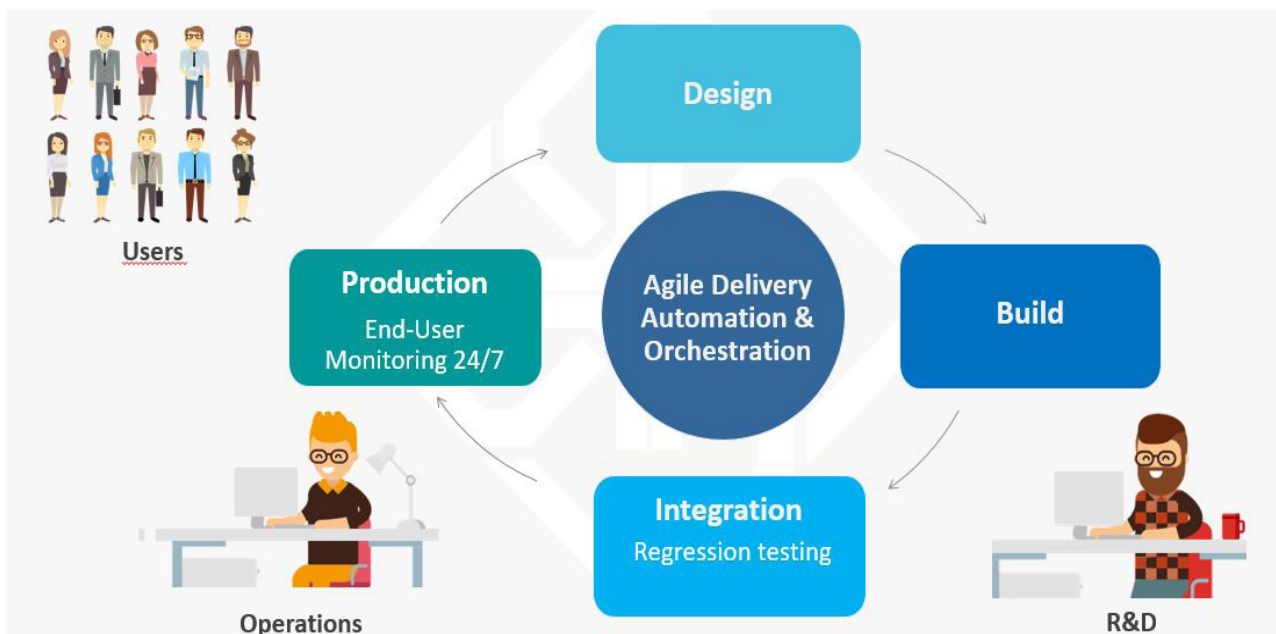
The performance of your business and of your business applications are entirely interrelated. Testing and monitoring your enterprise applications enables you to minimize the risks of unavailability of your critical applications and ensure a satisfactory experience for your users.

1. Causes of poor performance

Performance problems can arise at any time and any place along the application delivery chain (from the application provider all the way to the end user). Any number of causes may come into play:

- functional bugs that can be found in all environments
- bugs related to the environment and data
- intrinsic slowness
- slowness stemming from load (platform capacity)
- total or partial unavailability of functionality

These issues can be anticipated upstream in a continuous integration cycle like the one shown below:



Source : ip-label

According to Forrester Research, **90% of IT managers say their staff is engaged in a daily struggle against tech problems** that they have no way of detecting.

At the same time enterprises use APM tools to monitor **only about 5% of their critical applications**, and **most of them are monitored from the server side**.

It is vital **to monitor critical applications from the end-user's standpoint**, whether this means from a mobile phone, a tablet, a laptop or a desktop PC. This way, your IT teams can speedily detect, prioritize, and resolve problems to improve workforce productivity.

Below are described a variety of tools for **monitoring your critical applications at the pre-production and production stages**.

2. Pre-production testing: regression and load

It is important to detect and fix bugs or faulty performance as quickly as possible, to avoid expensive management of mass incidents in addition to costly 'firefighting' and deployment of associated patches or workarounds.

Regression testing

The point of **regression testing is to ensure software quality**. These tests are generally conducted in an integration environment as similar as possible to the production environment. Beyond unit tests (code functions), **end-to-end testing is needed, with complete transactions or user journeys, in order to fully validate operation**. Automating these tests is a key factor of success in this context where updates are so frequent.

It happens, nevertheless, that some bugs appear only under live operating conditions (load, environment, flows, etc.). This is one of the reasons why it is so important to monitor applications in production in addition to pre-production.

Load testing

Load testing pushes the platform (application, databases, etc.) **to the limit in order to identify the precise weaknesses and maximum capacity** of a given infrastructure budget.

The purpose of load testing is to validate technical choices and ensure that a new version or new platform can handle the expected amount of traffic. It also serves to clarify and visualize poor response times in order to determine the best fixes and then, afterwards, to check that the changes yield the expected results.

The best time to undertake load testing depends on your context and needs. For a new application/version, **load testing can be performed prior to production rollout or upon delivery of a new version**. For a change to the infrastructure, load tests can be run when migrating to the cloud or upon any change to the architecture.

The technical specifics of load testing differ according to whether thin client business applications (web applications) are being tested, or thick client applications.

To test the effects of load on a virtual desktop (Citrix VDI, by way of example), the console controls injectors, which instantiate Citrix sessions and run scenarios in the thin clients.

For thick client load testing, the console likewise controls the injectors which, in this case, are installed directly on the VDI server. The injector opens one user session per virtual client in which it runs a scenario.

The load test can record a number of metrics, including:

- **business metrics** (number of unique visitors per hour, number of operators, etc.)
- **client-side technical metrics** (availability, response times, bandwidth usage)
- **server-side systems metrics** (CPU usage, RAM, disk I/O, network interface, etc.)

These metrics make it possible to determine the maximum capacity of the architecture and its sensitivity to load on user journeys.



Case study VDI load testing

A French international transport company decided to rationalize IT for all of its locations in France. It called on a technology partner to set up a Citrix/VDI infrastructure to provide the transport company's workforce with a virtual desktop delivering applications such as Lotus, Office, IE, and so forth.

*To validate the technology partner's service, **load testing** was carried out by a third-party specialist simulating a load of 300 users to:*

- *Validate the infrastructure*
- *Define KPIs and SLA thresholds*



The benefits

- *Definition of 3 KPIs on the basis of test results*

1. *Time to open a Citrix session*
2. *Time to open a Citrix-delivered application*
3. *Response time of applications under load*



Load testing resulted in the validation of the infrastructure's adequacy to handle a specified load of user traffic and the definition of realistic SLA thresholds. In this way, IT services were set up based on clear, objective figures that were agreed on by both the transport company and its technology partner.

Load testing services (thick and/or thin client) make it possible to reproduce and identify anomalies so that you can accurately size your infrastructure.

3. Monitoring business applications in production

After the application has been rolled out to production, it needs to be checked and monitored **in operation 24/7 to make sure it is running correctly at all times.**

Several types of event may affect the performance of enterprise applications in production:

- defective hardware (disk, server, etc.)
- changes (configuration, patches, updates, etc.)
- latent intermittent malfunctions related to traffic or events in production
- the user's context: workstation, access network, etc.

Monitoring applications from the end-user's standpoint enables you to speed up incident resolution by detecting problems early on, **substantiate user complaints, generate KPIs, and better understand the root causes of performance issues.**

Substantiating user complaints

Because user experience is naturally highly subjective, a way of justifying or substantiating user complaints with objective measurements is needed.

Without tools that show you the end user's view, none of the parties involved in application delivery has insight into the many possible root causes located from one end to the other of complex delivery chains. The examples below illustrate what this means in actual situations:



From the IT department's standpoint

In general, IT departments have several vectors of information about the performance of their digital tools: user feedback via the helpdesk, network or server performance metrics and indicators, etc.

If the enterprise calls on a provider (outsourcing), it wants to make sure that this technology specialist is providing quality service. The enterprise needs to be able to talk with its provider in the event of an IT crisis or incident. Enterprise IT has no way of knowing whether the problem comes from the provider or elsewhere.






From the provider's standpoint

Because the application delivery chain is so complex, it is a major challenge to proactively detect service disruptions that affect its customer's end users. From the provider's standpoint, everything may seem to be working just fine... The provider is unable to determine exactly where the problem is occurring if it happens outside of its own scope.



From the end-user's standpoint

Problems may arise at any point in the delivery chain, but end-users shouldn't have to know or care about that. In fact, when they contact their helpdesk about performance problems, IT indicators often show that all is well. This apparent contradiction only adds to the user's confusion and wastes more time.



A number of solutions are available for monitoring the performance of critical applications in production. This white paper discusses 4 approaches:


- **synthetic monitoring,**
- **passive monitoring,**
- **troubleshooting (APM),**
- **network analysis.**

These solutions work alongside each other to shed light on every aspect of the delivery chain for comprehensive insights into performance deficiencies.

Monitoring with robots (synthetic monitoring)

The principle of **end-user synthetic** or '**active**' monitoring consists of **interacting with applications as a human user would**. This is done **by running scripted scenarios at programmable intervals** to simulate user transactions (or journeys), in other words, interactions between the user and the application.

Only end-to-end measurement (user workstation <==> network <==> data center), from the user's point of view, can offer insights into the quality perceived by users:



Factual measurements are made by **robots** which simulate the same kinds of transactions that end-users would perform. As each successive step of the transaction runs, the robot measures the application's **response times** and **availability**. The resulting metrics indicate the time it takes for the application to respond, breaking down each of the steps involved from end to end to provide precious diagnostic information (screenshots, waterfall of HTTP and network loading times, etc.).

This type of monitoring runs scripted user journeys around the clock to immediately detect any performance problem that might affect end users. These measurements provide real-time insights into the **availability, response times, and performance of the enterprise's critical application services**.

With a solution that uses robots for monitoring, the enterprise is alerted of any degraded performance for **any type of application deployment** (monitoring of thick-client or mobile applications, web browsers, RDS clients, Citrix clients, VDI, etc.). It can also **anticipate user complaints** to handle them with greater efficiency. The solution also helps to shorten the time it takes to detect and resolve incidents. Metrics furthermore facilitate triage of incidents according to their business priority (impact on the user).



Case study ERP and CRM applications

A European regional energy company with enterprise applications (ERP, CRM), Office software suite, Citrix platform, and business-specific applications decided to set up monitoring on the user side to measure application service levels in real time.

After its APM specialist set up robots at two locations, the energy company was able to observe differences in application performance between the two sites. Location-related problems were clearly identified through analysis of the response time and availability metrics supplied by the robots.



Case study PLM applications

An aeronautics manufacturer uses the Windchill product lifecycle management software (PLM) to manage inventory. In addition to overseeing applications, the company's application operations unit also handles incidents and user complaints.

After their APM provider set up a synthetic monitoring robot to conduct typical transactions on the application in question, the operations unit was able to collect the metrics it needed to document response times in continuous 24/7 operation, from the locations of the users and the various project participants (engineering teams, production teams, and customer service) distributed across Europe.



Case study Outsourcing and external metrics

An organism responsible for validating pharmaceutical products relies on Citrix-delivered business applications for its workforce. The in-house technical team of 10 was not enough to cover every need, and so it was decided to partially outsource IT requirements. The technical team is therefore dependent on the quality of service supplied by its technology provider. Internal users complained to them that the applications they were accessing via Citrix were slow.

But was the problem located in internal IT or rather on the technology provider's side?

To gather conclusive data to show their provider, IT management decided to call on an external specialist to monitor the Citrix-delivered applications. With the metrics resulting from targeted monitoring of the performance of the critical business applications, the IT department was able to demonstrate to the technology provider that service was objectively slow, and that the quality of service fell short of the specifications in the SLA.

RUM monitoring (passive monitoring)

Unlike monitoring with robots, **RUM (Real-User Monitoring)** observes real users, capturing comprehensive data in real time about the level of performance they are experiencing.

Passive monitoring offers a **precise view of performance in every usage context**. A valuable complement to active measurement analysis services, RUM solutions keep application managers and operations staff informed of the performance experienced by 100% of real users.



Measurements are conducted on the user's device by means of a component (JavaScript) embedded in the application. It collects the response times of each critical step and compares them according to user location, hardware characteristics (PC, smartphone, tablet, etc.) and software (OS, browser and browser version, etc.). This type of monitoring can also correlate performance with the number of users to identify any deterioration due to traffic surges.

With a passive monitoring solution, the ROI is fast because you can immediately visualize performance and usage results upon rollout to production or after optimizations to a website in production, without any other test instrumentation.

Troubleshooting (APM)



Whether the need is to measure and improve end-user experience or to lower the operating costs related to incident management, **troubleshooting agents installed on servers can monitor 100% of transactions and code.** They offer full visibility into the internal operation of applications, down to the line of code.

With a troubleshooting solution, the enterprise can monitor legacy applications and complex environments (Java, PHP, Microsoft.NET, C/C++, etc.). All metrics are searchable via SNMP and SOAP web service call.

Monitoring is carried out from data centers. It yields targeted diagnostics including **error codes, exceptions, view of the code responsible for slow response times, etc.** It can also identify the transactions which consume the most resources and the slowest SQL requests.

Nevertheless, when there is no user activity, there is no way to detect malfunctions. This can be the case at night time, or when network access is disrupted.

The cost of the intrusive agents and their impact on overhead may also be a disadvantage.

Network analysis

To determine whether performance shortcomings or deteriorations noticed by users originate in the network calls for a diagnostic of network performance.



Network analysis offers a wide-angle view of application performance by measuring 100% of transactions and 100% of users end to end via the network. It provides metrics that translate network bottlenecks, the impact and root causes of incidents (for example, resource-hogging applications, bandwidth usage, etc.), and more.

No specific configuration is required for network analysis. A virtual or physical network capture point is installed non-intrusively (no agent, no changes to code) to capture all application transactions. Such solutions monitor IP connections, TCP retransmissions, server response times, etc.

This approach is limited in terms of application analysis of malfunctions from a transaction point of view, and when traffic is encrypted.

To gain comprehensive insights into application performance, it can be useful to combine different types of solution to complete the picture from all angles: robots to corroborate user experience, APM troubleshooting agent to analyze application defects, and network analysis to identify low-level contention.

Conclusion

This white paper demonstrates **how crucial it is for enterprises to implement monitoring tools and full metrology of the entire delivery chain of mission-critical applications**. But which solution is the right one for a given enterprise?

The choice of an application monitoring solution cannot be made without a full analysis of your needs, as well as an in-depth analysis of the features proposed by APM providers. A good provider must be able to accurately define the customer's needs... in consultation with the customer.

Dialogue is the key. This implies a human relationship between the two parties, because understanding expectations and formulating the best response means listening to the customer.

Among the criteria that must be set forth clearly, are of course, the type of application to monitor (thick-client and/or thin-client, Citrix or web, etc.). The solution proposed must also be able to interface non-intrusively with the customer's information system, including upstream (pre-production) in DevOps mode.

It is up to the application monitoring specialist not only to provide solutions to measure the performance of business applications, but above all to advise customers, work closely with them, and even take care of certain tasks on their behalf.

With this in mind, the APM provider may propose additional approaches to benefit the client enterprise in response to its specific constraints. For example, if an enterprise does not have the necessary staff to manage the type of project that is drawn up, the provider may propose a SaaS type of solution or an installation that can be managed remotely.

Furthermore, since the provider's role includes advice, assistance, and guidance, it must anticipate and arrange to be able to intervene remotely on the monitoring solution, whether it is on premise or SaaS. In all cases, the provider must be available to the client enterprise to present results, explain them, and recommend optimizations.

To conclude, monitoring the performance of business applications is a technical service performed by competent, experienced specialists. But it is also a service characterized by the human dimension from the outset, with personal attention to defining the project, setting it up, throughout application monitoring, and in making the results meaningful to the customer.

About the ip-label group

ip-label has been assisting enterprises in their digital transformation for over fifteen years. A leading specialist of user quality of experience, the ip-label group offers an extensive range of APM (application performance management) solutions for analyzing and measuring the performance of all digital services including web, business apps, mobile, telephony, video and voice.

Want to know more about application
performance management?

[Request a demo](#)



www.ip-label.com

+33(0)1 77 49 53 00

marketing@ip-label.com



Paris | Madrid | Stockholm | Helsinki | Shanghai | Tunis

©ip-label 2019. All rights reserved