

IN PURSUIT OF P.R.A.I.S.E:
DELIVERING ON THE SERVICE PROPOSITION
What to Look For When Considering On-Demand Applications



Christine Mason, CEO

SpringCM

Best Practices in Operations, Performance and Security

Table of Contents

Seeking P.R.A.I.S.E: Best Practices in SaaS Operations	2
Performance: Measurement, metrics and monitoring	3
Reliability: Components that ensure optimal user experience	5
Availability: Delivering on the “five nines”	6
Information Stewardship: Information protection short and long term	7
Scalability: Ensuring a sound capacity management program	9
Enterprise Dependability: Evidence to support confidence in business growth	11
References	12

Seeking P.R.A.I.S.E: Best Practices in SaaS Operations

You've made the decision: the overall ease and economics of implementing Software-as-a-Service (SaaS) appeal to you, and you have selected an application that delivers what you need for your business.

The last hurdle is the service component itself — making sure that critical factors like security, speed, and reliability that will affect your users' experience meet your standards. When you install or deploy software internally, your own Information Technology team handles operational items like these. With SaaS, where you are outsourcing operational processes, you need to make sure your vendor has the right approach for providing the service you deserve.

The good news: SaaS solutions are for the most part built to deliver excellent security and performance levels: making applications run well from a central, secure point is the core business of a SaaS provider. Large investments to ensure service levels can be shared across many customers, and often create better operational results than systems deployed in house. Nonetheless, not all SaaS providers are created equal.



What makes a good SaaS partner?



Built for better security and performance



Culture of service, good processes and practices

Delivering on a Service Proposition

A service proposition begins with intention: being tuned into the needs of the customer. Software companies that started as “product” companies — where you build, package and ship a product — have a different perspective, and have aligned their operations differently. Being a “service” provider involves more than hosting applications. Before you get to best practices, metrics and measurements in operations, ask yourself: Does the company seem to have a service mentality? Do they see themselves as stewards of your information or just operators of infrastructure? How can you tell?

When a company combines a powerful service culture with good processes and practices, you have found your SaaS partner.

Look at the language the company uses about service, and consider how well it listens to your needs and understands the value you seek to create in your environment. Do they offer live sales and support people to assist with questions? How earnestly and directly are issues handled? Does the company offer service level commitments and publish its performance statistics for customers to see? With that in mind, here are some guidelines that a thoughtful consumer can follow when verifying a potential solution's operational effectiveness. We call it seeking “**P.R.A.I.S.E.**”

Performance:	Measurement, metrics and monitoring that raise the bar
Reliability:	Problem submission, determination, resolution and tracking
Availability:	Service commitments for when you need it most
Information Stewardship:	Obligation to protect your information now and in for the future
Scalability:	Capacity programs encompassing a 360° view
Enterprise Dependability:	Customer growth, financial strength and value of SaaS

Performance

When assessing performance, what you are really asking for is a guarantee that the SaaS provider will deliver its service consistently and seamlessly, responding quickly to the “clicks” of your users. Speed matters: studies show that, on average, a user expects a response in no more than 2 seconds, is willing to wait up to 5 seconds, and has almost always given up by the time 8 to 10 seconds has elapsed.

Many companies measure server performance alone, or the speed at which the back end server responds to requests. This does not matter much to end users, however, unless that speed translates into a rapid response to their requests. Vendors also boast of how many transactions per second (TPS) the system will handle — but again, TPS in itself does not necessarily translate into a rapid response to the user. Other factors like the size of the page being returned and the complexity of the logic in that page will also have a significant effect on performance, as do bandwidth and the size of the data set being processed.

Here are a few more specific performance-related questions to ask:

WHAT TO ASK	PERFORMANCE: WHAT TO LISTEN FOR
<p>How do you measure response time?</p>	<p>A vendor with a service mentality will put themselves in the shoes of the user. They will use external measurement tools to look at both server response time (how quickly the application does its work), and browser response time (how quickly the page finishes loading from the end user’s viewpoint) — and make sure to optimize both.</p> <p>It’s a fine point, but it pays to ask what the measurement points are. For example, is an end-to-end use case being measured, or just individual page response? In evaluating whether the numbers are meaningful to you, you may want to check to see that the data points are gathered in a way that matches the way you will use the system — for example, are the pages and processes involved in your most frequent transactions the ones that are being measured?</p>
<p>How do you monitor performance?</p>	<p>The company should work with an external monitoring company, such as Keynote Systems, who measures performance data and trends over time, and provides quality, audited, independent information to the vendor to allow them to continuously improve the solution. Some application providers have chosen to do similar instrumentation of the application themselves.</p> <p>Best practices dictate that a SaaS provider should be monitored externally, and should be alerted in real-time when performance drops below a given threshold.</p>
<p>How quickly do your users see a response to their clicks?</p>	<p>Now that you have clarified how your potential SaaS provider monitors the performance of the system, ask for data points. We recommend that you look at daily data, since many reporting packages “smooth out” outliers over a long period of time. Nothing less than a month should be observed. In addition to median response time, you might want to ask about standard deviations — that is, the predictability and stability of actual responses, not just averages.</p>

WHAT TO ASK	PERFORMANCE: WHAT TO LISTEN FOR
<p>Do you have a service level commitment or agreement?</p>	<p>The SaaS provider should clearly outline what performance levels the customer can expect. Ideally, an overall response time for the application should be defined within percentiles:</p> <p>Currently, a good target for SaaS SLCs is to have 95% of transactions with a server response of less than 5 seconds and 99% of transactions less than 10 seconds. Exceptions depending on load are acceptable, as long as they are clearly identified (e.g., transactions whose response time is dependent upon load-file-upload transaction). There may be other acceptable dependencies based on outside providers, for which the SaaS vendor is not responsible; for example, the speed of the Internet backbone. This kind of exception in a service level commitment, while not ideal, is acceptable if identified up-front.</p> <p><i>Note: There is a difference between a service level commitment and an agreement. The former is a non-binding version of the latter. SLAs generally come at a premium price, because of the reporting and management requirements, such as regular reporting that summarizes the performance of the system for that reporting period by customer.</i></p>
<p>What are your escalation procedures for performance-related issues?</p>	<p>The vendor of the hosted application should provide a defined management escalation process that can be used when response time service level problems are not resolved in a timely manner.</p> <p>Ask your vendor to see their published escalation procedure, and look for the following best practices:</p> <ul style="list-style-type: none"> ▪ Clear definition and classification of issue types (e.g., severity criteria) ▪ Tiered escalation within an organization by issue classification ▪ Clear policies on when customers receive service notifications <p>Ask your vendor how many performance alerts need to occur before real-time action is taken, and look for responses like “ONE” to assess their sense of urgency around customer experience. Anomalies should not be treated as “we don’t know” or “it’s not worth looking into yet,” but should immediately enter a queue of discovery. The goal is to get to the heart of the issue quickly and see if it suggests any larger issues.</p> <p>While it is up to you to determine whether the performance measurements meet the needs of your business, it is safe to say that, as regular Internet users, business personnel expect the same response times on their SaaS application that they get on Google, eBay, or Amazon, for example.</p>

R.eliability

Availability and Reliability are intertwined, but they are not one and the same. In systems with subcomponents, certain functions may be unavailable in spite of overall uptime. The following questions will ferret out unexpected component issues that may affect your experience.

WHAT TO ASK	RELIABILITY: WHAT TO LISTEN FOR
<p>When the hosted application is available, do all the application components and functions work as intended?</p>	<p>If the vendor states that the application is available, but key functions are inoperable, the system is not fully available, even if this isn't immediately apparent to the user. Application reliability can be indicated by the number and types of application problems (see additional questions below).</p>
<p>Does the hosted application service provide a database of known problems?</p>	<p>Your SaaS vendor should provide a place for customers to review all known (currently outstanding) and past (resolved) problems. On a regular basis, the vendor should provide a statistical summary of problem management information.</p>
<p>Does the vendor provide a mechanism for problem submission, determination, and resolution (with an associated escalation procedure for problems which are not resolved within the standard timeframe)?</p>	<p>Best practices are to provide a standard mechanism for problem submission, determination, resolution, and tracking that keeps the customer informed as to the status of all their outstanding problems.</p>
<p>Does the vendor provide a definition of problem severities with associated resolution times?</p>	<p>As part of problem resolution and escalation, your SaaS provider should have an objective definition of problem severities along with associated resolution times for each severity.</p>
<p>Are there any known single points of failure in the system today?</p>	<p>The goal is to have redundancy for each component of the system; however, this is frequently not the case. If there are components of the system for which there is no redundant backup, what measures are now being taken, or planned for the future, to ensure that the entire system has a fail over plan?</p>



Availability and Reliability are not the same



The goal is to have redundancy at each point



Best practices ensure the promise of on-demand

A.vailability (uptime)

Availability helps you assess whether the system will be there for you “on-demand” as promised. The telecommunications companies speak of “five-9s” availability, meaning 99.999%. While “five-9s” is the extreme in availability, and virtually no modern software attains it, 99.9% availability is highly attainable in a mature and well-managed system. The percentage availability of a system over an entire year is a good indicator of how often you will be without service from one of your critical business systems — 99.9% availability translates to 9 hours of downtime a year.

WHAT TO ASK	AVAILABILITY: WHAT TO LISTEN FOR
<p>Does the hosted service include a defined service level commitment or agreement for system availability?</p>	<p>The service provider should clearly outline the service level that the customer can expect. This metric consists of two dimensions: network availability and application reliability.</p> <p>Network availability indicates whether or not the application was “accessible” from an external point within the network. Application reliability indicates whether or not all the application functions were operational (when the system was available). In doing so, any availability exceptions (e.g., standard downtime for application/system maintenance) should be explicitly identified.</p> <p>If a service level agreement is provided, it should contain monthly service level reports from an audited source. A service level commitment may provide only the company’s own internal data and reports.</p>
<p>Does the vendor provide a defined management escalation for availability related problems?</p>	<p>The vendor of the hosted application should provide a defined management escalation process that can be used when availability problems are not resolved in a timely manner.</p>

Note: Network availability is defined as the ability to access the system from ANY external point within the network. Because of the complex nature of the Internet backbone, problems with ISP availability can cause a given user to perceive the system as being “unavailable,” when in fact only their ISP or network backbone are affected.

I.nformation Stewardship

(data protection, both access and backup)

While information stewardship includes information security — keeping your information safe and locked down — this is only part of the picture. More broadly, you want to be sure that the vendor understands their deep responsibility for your information, and takes seriously their obligation to protect it for your current and future use.

WHAT TO ASK	INFORMATION STEWARDSHIP: WHAT TO LISTEN FOR
<p>How do I know that my data is secure?</p>	<p>Your SaaS provider must take on the role of "data caretaker." Within this context, data security extends across many dimensions, such as:</p> <ul style="list-style-type: none"> ▪ Physical security ▪ Application security ▪ Database security ▪ Network security ▪ User/account security <p>The hosted application service must address key data protection and security issues within each of these areas.</p>
<p>What is your policy regarding the communication of any security-related violations?</p>	<p>The vendor must embrace a policy of full disclosure in the event of a security-related violation.</p>
<p>What disaster recovery provisions have you made to ensure continued operation in the event of a disaster to your primary facility?</p>	<p>The vendor must have documented disaster recovery plans in place to ensure continuity of operations in the event of a disaster which affects its primary facility.</p>
<p>Who owns the data that is housed within the hosted application?</p>	<p>Unequivocally, the data is owned by the customer who is contracted to use the service. In its role as data caretaker, the vendor is granted limited rights to that data. In no way does that vendor own that data, or any derived data, unless explicitly agreed to by the customer.</p>
<p>How do you protect the perimeter? Do you employ Intrusion Detection Software (IDS) or Intrusion Prevention Software (IPS)?</p>	<p>Intrusion detection and prevention have been rising in importance in recent years. Malicious attacks on SaaS vendors are increasing, but so are the maturity of products that combat these threats. Because of the sensitive nature of these solutions, SaaS vendors may be reluctant to divulge their products and plans. They should, however, be able to clearly articulate the costs/benefits of these solutions.</p> <ul style="list-style-type: none"> ▪ Physical product placement (e.g. firewalls and server farms) ▪ Protection against <ul style="list-style-type: none"> ○ Unauthorized access ○ Malware ○ Denial of Service ○ Application/Database attacks <p>Many companies provide solutions for any need and budget, such as the products that fit within Cisco's Self Defending Network, or the open-source solution Snort.</p>

WHAT TO ASK	INFORMATION STEWARDSHIP: WHAT TO LISTEN FOR
<p>How often do you update your virus profiles (applicable to Windows only)?</p>	<p>A SaaS provider should have a plan in place to deal with the continual inoculations available to Windows viruses. Third party software such as McAfee or Sybase, for example, can be used to notify of updates and manage the installation.</p>
<p>Do you run regular security audits?</p>	<p>An external audit of a SaaS provider's security is an important exercise that should take place regularly regardless of provider size or maturity. These should be conducted at least annually, but preferably semi-annually, with an outside certifier of the data.</p>
<p>What are your policies for physical and remote access?</p>	<p>Any SaaS provider must control access to hardware, software, and data — both physically and remotely. In doing so, they should be able to articulate control mechanisms, policies, and procedures in a manner that's clear to a prospective customer. Physical security mechanisms and procedures must strictly verify the identity and authorization of any and all persons who will have physical access to the site. Effective approaches to physical security include key cards, mantraps, video surveillance, and biometrics.</p> <p>Virtual facilities operations can exceed the strongest physical infrastructure. Measures to look for in a qualified SaaS vendor include: close monitoring and logging of remote access to systems and data; strong ID/password security mechanisms (e.g., regular password change, limited user access, etc.); IP subnet security implemented within the site to ensure that remote access is only allowed from specific "authorized" subnets; and VPN technology (with strong encryption) to ensure the protection of any transmitted data.</p>
<p>How is access to the data audited?</p>	<p>Equally important as the ability to secure the data in the system is the ability to know who accessed it and when they did so. A full audit trail for each object in the system should be available indicating who accessed the object, operations performed, and when those operations took place.</p>
<p>How is data in the system secured on an individual or group level?</p>	<p>As in any large system, the provider should have established rules for who is allowed to see and modify data. At minimum, these rules should apply on both an individual and a group level. Optimally, data should also be able to be secured based on a user's role.</p>
<p>What are the various levels of security that can be assigned to a given object (document, record, etc) in the system?</p>	<p>At a minimum it should be possible to secure any object for read, read/write, delete, and edit permissions. Beyond these basic security measures, there are additional security levels that can or should be available based on the type of system and the applications the user plans to build.</p>
<p>How is the physical hardware for the system protected?</p>	<p>Obviously, access to the physical hardware upon which the system is running should be tightly restricted. It is useful to know the rules for granting access to the physical system and how individuals are monitored when they are working on the hardware. In a perfect world, no single individual is granted unsupervised access to the system.</p>

S.calability (planned for growth for concurrency and bandwidth)

Unlike installed applications, generally administrated directly by in-house personnel, the SaaS provider delivers applications to a wide variety of users who generally can self-provision new users, add content, etc., without enablement or intervention from the SaaS provider. For this reason, the vendor needs to be adept at predicting peak loads on the system, and planning for scalability. This will ensure that you can increase your utilization as promised, on-demand.

WHAT TO ASK	SCALABILITY: WHAT TO LISTEN FOR
<p>How many concurrent users is the system designed to support?</p>	<p>User concurrency is a simple metric often used to represent the "load" on a system. This metric is typically correlated with the response time service level over time to demonstrate the ability of the system to maintain its performance as the user load increases.</p> <p>A SaaS provider should be able to provide any prospective customer with an analysis that illustrates this correlation over time (i.e., response time service level vs. concurrent users). This analysis allows the prospective customer to assess two critical aspects of this provider: 1) the extent to which the user load has increased over time; and 2) the ability of the provider to meet its response time service level commitments in spite of the increased user load.</p> <p>When analyzing this, a prospective buyer should make sure that these measurements relate to "peak hours," and are not averaged across an entire day.</p>
<p>Are there known bottlenecks in the scalability of the system today?</p>	<p>As a customer, you are really asking if you can expect any challenges as a user, and trying to make sure that the vendor is prepared to provide excellent service both now and in the future.</p> <p>Nearly all well-known systems have identified scalability bottlenecks. The question then becomes, what proactive steps are being taken to eliminate these bottlenecks as the system grows?</p> <p>A good vendor will have a good handle on each component of the system, as well as its pressure points, and will make sure capacity is built in to handle customers' peak demand. Things to look for in a qualified vendor: a capacity plan, clearly defined alerts, and triggers for expanding capacity, and clearly understood costs and timelines for scaling the operations.</p>

WHAT TO ASK	SCALABILITY: WHAT TO LISTEN FOR
<p>As your service grows, how will you ensure that my response time and availability will not be affected?</p>	<p>In keeping with its response time and availability service level commitments, the vendor of a hosted application must implement a formal Capacity Management program to ensure that the system capacity is increased in order to meet increased demand from new and/or existing users.</p> <p>A formal capacity plan from the service provider is in effect a trade secret, and it is unlikely that it will be shared outside of the company, as it would allow an external company to re-engineer the provider's growth goals. As a customer, what you are looking for here is a discussion of how well the vendor is anticipating its future capacity needs, and making sure they will have enough data to see that your needs are met.</p>
<p>How far into the future does your Capacity Management program go?</p>	<p>Many vendors take a reactive approach to system capacity and only increase it in response to customer complaints about response time. A good vendor will have a Capacity Management program, which provides for increases in system capacity before current users are affected.</p>
<p>How do you ensure that your system will be able to perform with an ever-increasing number of users and their associated data?</p>	<p>In other words, "How do I know that the system will continue to perform adequately as its popularity and usage grow?" The SaaS provider's response should address two key points:</p> <p>1) Most large systems that are designed to handle extremely large loads incorporate two key architecture principles: distributed processing and horizontal scalability. Distributed processing allows the processing load to be spread across many smaller processing units within the system (this also provides redundancy, improving system availability).</p> <p>Horizontal scalability ensures that the overall processing capacity of the entire system can be increased by merely adding more "processing units" without needing to increase the processing capacity of an individual processing unit.</p> <p>2) The SaaS provider should be able to provide historical reports that demonstrate the extent to which they have tracked and monitored the system in the past, analyzed usage and growth trends, and increased system capacity to meet increased demand (while still maintaining their committed response time service level).</p>
<p>How much capacity do you keep freely available for peak-load anomalies? How do you track/measure load?</p>	<p>A provider should not be running at more than 80% utilization during their peak hours. In addition, they should be able to provide a prospective buyer with historical reports that illustrate the extent to which they have tracked and monitored system utilization (and its correlation with response time service level).</p>

E. nterprise Dependability

(vendor history, qualifications, backing)

Most purchasing departments have processes for assessing and mitigating the risk of any vendor. Specific to SaaS, as a business process owner, you should mostly be concerned about your data, and trying to make sure that your process will run in an uninterrupted manner.

WHAT TO ASK	ENTERPRISE DEPENDABILITY: WHAT TO LISTEN FOR
<p>How long have you been in business?</p>	<p>The Software-as-a-Service business model is relatively new. While you may find application service providers (ASPs) that have been hosting applications for many years, many of these vendors are merely hosting applications that have been developed by other companies.</p> <p>Software-as-a-Service represents the new breed of hosted applications, in which the provider is both the software vendor and the operations vendor. Not only does this vendor know how best to run your application, but they have the capability of enhancing and extending the application to better meet your needs. In addition, all customers are running on the same version of the core code base, which creates significant financial advantages for the customer.</p> <p>Look for evidence of customer growth, financial backing and strength, and SaaS rather than ASP economics.</p>
<p>What happens to my data if your business fundamentally changes in some way?</p>	<p>Business environments change constantly. If for any reason you need your data to be returned to you, any qualified vendor must have provisions for delivering it in a format that is usable to your business.</p>
<p>Can I have references from current customers?</p>	<p>Contacting existing customers can be a great way to gauge the current and past performance of the vendor.</p>

Software-as-a-Service combines the best of application design, development, and deployment, with the ease of provisioning and use inherent in consumer software. Along with this comes the expectation that the service level provided will perform at the same response levels as e-commerce and Internet content sites, and provide business grade dependability. When a company delivers on both sides of the SaaS promise — the software’s quality, and the overall service level — you have found a solution provider than can deliver for you.

References

About the Author



Christine Mason is co-founder and CEO of SpringCM. Christine previously served as CEO of Open Road Technologies (now Intellect), a venture-backed advanced semantics search technology company. Prior to Intellect, she was managing director of ClickMarkets, an early stage venture accelerator, whose companies include GoBeam, Mobitrac, LeadAmerica, and Holisticare, among others. Christine specializes in opening new markets and inventing business lines- she worked for many years in mergers and acquisitions, post-merger integration, and new channel and market growth with global companies such as Price Waterhouse Coopers, and Kraft General Foods, among others. Christine is on the Board of MSDSONline, Inc., and serves on the Software Board of the Software and Information Industry Association. She is a member of the Economic Club of Chicago. She holds a BA from Northwestern University, and a MBA from Northwestern's J.L. Kellogg Graduate School of Management. She speaks German and French, and is a frequent speaker and writer on technology, new venture creation and leadership. Christine and her husband John live in Sonoma, CA and Chicago, IL with their 6 children, 2 dogs and 5 cats creating joyful chaos (perhaps even more than life in the technology sector does). Together, they volunteer on issues related to clean power and global collaboration.

About SpringCM

SpringCM is the leader in on-demand document and content management, helping companies of all sizes, both domestically and internationally, automate their document processes, improve communication, and collaborate, both inside and outside their company walls. SpringCM meets the real need of businesses to manage the explosion of documents and other content types – from capture, storage and search, through collaboration, workflow, and document process automation. Today, customers such as Avon Products, General Electric, and Cox Communications are utilizing SpringCM's unique solution to solve their document management needs. SpringCM maintains offices in California and Chicago. The company is backed by Foundation Capital of Menlo Park, CA. For more information visit www.springcm.com or call us at 312.881.2026.