RingCentral

Understanding reliability

Understanding reliability

It's easy to confuse the concept of the reliability of a cloud communications solution with the related—but different—concept of availability. Let's try to simplify it.

Let's start with availability. Cloud software—despite many advances—isn't quite perfect. Solution providers sometimes need to take solutions offline for short periods to do maintenance (sometimes known as a scheduled outage). And sometimes circumstances outside the provider's span of control can take a system down (natural disasters, for example).

To accommodate for these things, providers give service-level agreements (SLAs) that guarantee as little downtime as possible. An SLA of 99% annually translates to 7.2 hours of downtime a month, for example. On the other hand, an SLA of 99.999% or "five nines" (often viewed as the gold standard), translates to only 26 seconds of downtime per month.

The reliability of a cloud solution, on the other hand, comes down to the probability that the solution will deliver what you need it to over a period of time. It's really a reflection of how well the solution will perform under real-world scenarios—things like whether it'll deliver quality service as you grow, how it will handle network issues and other system failures, and the degree to which it'll secure your data.

Reliability consists of four key components: scalability, redundancy, quality of service, and security. The checklists on the following pages should give you an idea of the important questions you'll need to ask to make sure you're selecting a reliable cloud communications solution.

Evaluating scalability

Scalability is one of the key reasons for the popularity of cloud computing, particularly its ability to provide on-demand capacity with minimal lead time and effort. Of course, that all depends on your cloud communications provider. You'll need to look for cloud providers with well-established histories of high-volume, high-quality service (number of customers, number of voice minutes, etc.). Use the following checklist to assess your cloud communications provider.

What to ask: How does the cloud service provider manage their system capacity both for organic growth and unexpected demand surge?

What to look for: Ideally, you want a provider who actually has no limits. Practically speaking, your provider must have solid instrumentation to measure their capacity, a strong process to manage it proactively, and, most importantly, an architecture that supports a seamless elastic model of expansion. In order to do that, look for providers that feature a modular architecture, such as POD architecture. With this kind of robust architecture, there's no limit to what the system can handle in terms of users or traffic. If a single instance is full, another instance can be built instantly, and the accounts may be federated together.

This type of capacity management scheme proactively builds capacity when added demand is anticipated, preparing the system for any increase in traffic. Data center virtualization and a fully automated build-and-deploy engine are two excellent ways to do this quickly and seamlessly.

What to ask: Can the provider offer any details around its capacity for number of calls and calls per second, connectivity, long distance, local, etc.? What are the service provider's key control indicators (KCI)?

What to look for: Phone capacity is often gauged in two metrics: the percentage of a customer's current capacity that can be added immediately, and the number of minutes the provider manages annually. Ideally, look for providers who can handle 2x your capacity immediately, if necessary. As far as minutes are concerned, the larger providers handle more than 10 billion minutes of voice traffic per year.

In addition, if you do business internationally, it'll be important to find a provider with a large global reach that offers native public switched telephone network services (the circuit-based phone system over which the world's landline calls are made), as well as service availability in all major geographies.

What to ask: How easy is it to add or delete a large number of users from the system?

What to look for: A cloud communications solution should be easy to scale, with options to add users, telephone numbers, and extensions in real time. These should be available instantly with zero downtime and without otherwise affecting other users.

Look for an easy-to-use admin portal that allows a system administrator to control every aspect of the solution, including moves, adds, changes, and deletions (known as MACDs). In today's world, it's also critical that your admin can fully control and monitor your solution from anywhere at any time, including via both web and mobile applications.

Evaluating redundancy

Like it or not, the “9-to-5” workday is a relic. Today, work happens around the clock. In large part, that’s due to the staggering amount of work conducted via mobile devices. Regardless of the cause, this new work paradigm requires reliable systems 24/7. Employees—and customers—expect nothing less. So network disruptions simply can’t be tolerated—whether it’s a system failure, a natural disaster, a hack, or just human error. The costs can be devastating for businesses. When evaluating redundancy for a cloud communications provider, consider the following components.

What to ask: What tier of data centers does the provider host?

What to look for: Your provider should have data centers in multiple geographic regions (geo-redundancy). Tier-4 certification is currently the highest classification among data center facilities. Such data centers are the most reliable because they’re fully redundant and fault-tolerant, and no single outage or error can shut down the system. They have redundancies for every process and data protection stream, so data continues to flow regardless of day-to-day maintenance of systems. Each features a redundant power supply.

In terms of service architecture, active-active network design provides for maximum reliability and fast recovery. Overall, you’ll want a provider that guarantees 99.999% system availability and uptime for its service, often viewed as the gold standard.

What to ask: How often does the provider test its disaster recovery plan?

What to look for: Look for providers that have documented plans and policies and execute failover and data recovery plans at least once a year.

What to ask: How does the service provider ensure redundancy?

What to look for: As we discussed earlier, geographically distributed redundancy is critical in this situation. Primary and backup locations need to remain online simultaneously, with a primary instance in active mode

and a secondary instance in standby mode. If a system failure within one data center is detected, the redundant system—whether within that same data center or at another data center—should take over operations in accordance with internal failover policies and procedures. You'll want to make sure that your data, including phone service, service configurations, and messages, is fully replicated between locations in real time, with failover being built into the service. In the event a geographic disaster causes a data center failure, be sure that another major data center assumes immediate and complete system operations with no loss of functionality or customer data.

What to ask: Does the provider offer the following capabilities?

What to look for:

- The ability to add new features and test them without disruption of service
- The ability to upgrade users to new features with no loss of service
- Frequent rolling software upgrades as a means of getting the latest features immediately—and automatically

What to ask: What if my internet goes down?

What to look for: One of the consistent concerns about cloud solutions is that the dependence on the internet leaves companies vulnerable in the event of internet outages. And when we're talking about something as critical as communication and collaboration, those outages can be devastating.

Look for providers with solutions that include a means to deliver consistent service despite internet outages. Your provider should be able to deliver at least certain essential services in these situations, including:

- Emergency calls
- Outbound and inbound calls
- Extension-to-extension dialing

Your communications applications should run on both Wi-Fi and 3G/4G/5G for seamless communication in multiple network environments. You should be capable of making calls using a mobile phone's native cellular capability, while still maintaining the user's business caller ID.

Evaluating quality of service

The advantages of cloud communications and collaboration—cost, flexibility, immediate access to the latest technology—have become even more obvious in a world where remote work reigns. All those advantages, however, would prove meaningless if the quality of communication in the cloud suffered, particularly voice. Because voice data travels along the same internet "pipes" as other data in a cloud system, potentially creating bottlenecks that could hamper service, quality of service (QoS) from cloud providers has taken center stage. QoS is a set of parameters that allows you to classify and prioritize traffic on your network so quality won't suffer. Here are a few questions to ask about this important topic.

What to ask: Does the provider support internal diagnostics—the ability for internal self-detection, diagnosis, and reporting?

What to look for: Advanced QoS analytics are critical here. Look for analytics that track each leg of each individual call and provide call quality statistics. These help you diagnose and troubleshoot issues that may exist on the network, as well as identify what could be causing the issue. Your QoS dashboard should report on the mean opinion score (MOS), jitter, and latency.

You'll also want to look for:

- **Proactive monitoring** to identify potential issues based on patterns of degradation in call quality and to correct quality issues before they become disruptive to your organization.
- **Reactive problem resolution** to handle individual user escalations, find problematic calls and meetings, and identify their root cause.

What to ask: How does the provider ensure communications and transparency around service availability?

What to look for: Ideally, you'll want to work with a provider that maintains an updated site for customers to see all known system issues that could impact their service, preferably with drill-down capabilities to see the status of the network node where your account resides. This service portal should help confirm

whether there are outages, that endpoints are configured properly, and that environments meet network requirements and recommendations.

What to ask: What is the provider's methodology for assuring quality of service?

What to look for: Among the most important components of a strong QoS methodology are:

- Advanced packet-loss concealment technologies with call-quality algorithms to remove echo, background noise, and packet jitter.
- Direct peering with major domestic and international carriers to reduce network hop length and manage network performance across both PSTN and data interconnects.
- A global network for layer 2 data services, as well as a PSTN gateway infrastructure that places network performance directly under the control of the network operations center (NOC).
- Professional services and a team of experts to help organizations properly plan for bandwidth at each location.

Evaluating security

Of the four components of cloud reliability discussed in this playbook, security often gets the lion's share of attention—and for good reason. With the exception of your employees, data has become your company's most precious asset. And the prospect of that data getting compromised carries serious implications for a company's brand integrity, as well as its balance sheet. So when it comes to protecting your data, cloud providers need to deliver multiple layers of security—from an enterprise standpoint, but also specific layers for business process, applications, data, hosts, networks, and even physical security. When evaluating solutions, cloud providers should demonstrate a robust security program that includes policies and procedures around change management, access management, vulnerability management, incident response, fraud monitoring, audits, access reviews, training, and third-party testing. Here are a few questions to consider:

What to ask: What security certifications does the provider have? Does the provider help meet your regional or vertical compliance needs? Is the vendor willing to share documentation and details?

What to look for: There are a number of critical security certifications you'll want to be sure your provider has achieved or supports, including:

- SOC 2 Type II and SOC 3
- HIPAA
- HITRUST
- FINRA
- ISO 27001, 27017, 27018
- C5
- United Kingdom National Cyber Security Centre's (NCSC) Cyber Essentials Plus
- McAfee Enterprise-Ready Cloud Services (previously known as SkyHigh Enterprise-Ready)

What to ask: How is the provider securing the user interface?

What to look for: One key to securing the interface is supporting single sign-on (SSO), which provides better security with a central authentication point, limiting the possibility of phishing. Does your provider support it? SSO allows your employees to access all your applications with one set of credentials, which can include email, phone number, or username, along with the password.

What to ask: How is the provider handling robocalls?

What to look for: Robocalling has become particularly challenging recently. Automated calls either trying to sell products or entice people to provide personal data continue to increase. Your provider needs to help you mitigate this issue. Look for a provider that offers an intelligent spam detection and blocking solution, and is also actively participating in and gearing to support the new STIR/SHAKEN framework. When fully deployed in 2021, STIR/SHAKEN allows carriers to work together to identify calls where the caller ID doesn't match where the call originates. Providers using this framework can automatically block such calls. You'll also want to find a provider that uses AI and analytics based on blacklists of suspected robocallers, giving your users the chance to ignore the call or block it completely.

What to ask: Does the provider do penetration testing and regular risk assessment (by a reputable third party)? How often?

What to look for: You'll want a provider that performs a penetration test (internal and external) at least once a year.

About RingCentral

RingCentral, Inc. (NYSE: RNG) is a leading provider of business cloud communications and contact center solutions based on its powerful Message Video Phone™ (MVP™) global platform. More flexible and cost effective than legacy on-premises PBX and video conferencing systems that it replaces, RingCentral empowers modern mobile and distributed workforces to communicate, collaborate, and connect via any mode, any device, and any location. RingCentral offers three key products in its portfolio including RingCentral MVP™, a unified communications as a service (UCaaS) platform including team messaging, video meetings, and a cloud phone system; RingCentral Video®, the company's video meetings solution with team messaging that enables Smart Video Meetings™; and RingCentral cloud Contact Center solutions. RingCentral's open platform integrates with leading third-party business applications and enables customers to easily customize business workflows. RingCentral is headquartered in Belmont, California, and has offices around the world.

For more information, please contact a sales representative.
Visit ringcentral.com or call 877-596-2939.

RingCentral

776231624 10/2021

RingCentral, Inc. 20 Davis Drive, Belmont, CA 94002. ringcentral.com