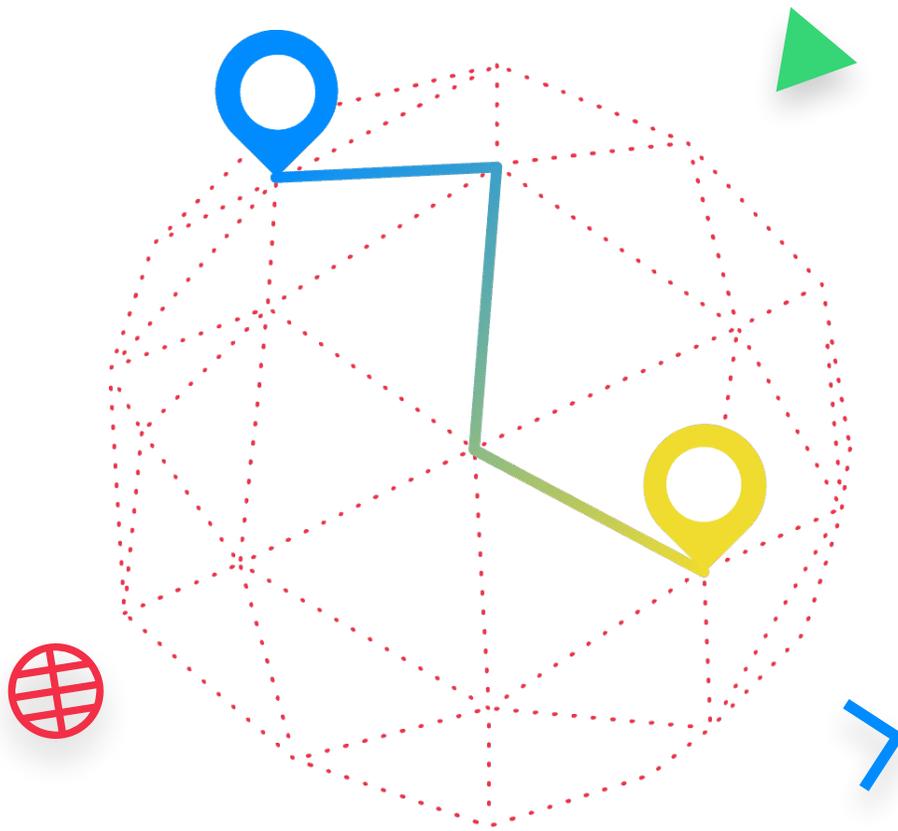

E-BOOK

Twilio's Global Infrastructure





Global infrastructure overview

Twilio's infrastructure is built to scale from the ground up, without sacrificing quality or security.

Twilio's Global Infrastructure provides resilience, redundancy, and low latency through Regions, Edge Locations, and private network connectivity endpoints. We continually evolve our offerings to increase the capabilities, footprint, and product availability we offer globally.

In the past year, Twilio has expanded the number of Edge Locations through which customer applications can connect to Twilio. We now offer Edge Locations on 6 continents, including 8 public Edge Locations with AWS and 7 private hardware connections available to customers through our Interconnect product. All customers can leverage private connectivity through AWS' backbone by leveraging Edge Locations for supported workflows.

Edge Locations allow customers to bring Twilio's public and private network connectivity closer to their applications for improved performance, security and reliability. Edge Locations can also process media handling nearer to the customers' locations to improve voice and video quality. Additionally, Edge Locations enable enterprise architects and developers to maintain their availability SLAs by selecting the Edge Location nearest their own infrastructure to provide the best performance for mission critical operations.

What's inside

- [Global infrastructure overview](#)
- [Global infrastructure components](#)
- [Software resiliency & testing standards](#)



Edge locations





Global infrastructure components

There are three main components of our Twilio's Infrastructure: AWS network backbone, Twilio's cloud and physical infrastructure, and how we meet our customers' need for direct, secure access to this infrastructure.

AWS

Twilio is resilient to both server and zone outages. Twilio operates in 8 AWS regions across multiple (3) Availability Zones per region. Availability Zones are physically separated datacenters in low risk flood plains, with uninterruptible power supplies that are fed from different electric grids from independent utilities, and onsite backup generation facilities. Twilio's services span over 10,000 EC2 instances to ensure redundancy and adequate capacity for every layer of our technology stack.

Twilio utilizes the AWS congestion-free backbone for inter-region network traffic to deliver low latency and high quality of service media and signaling to our customers. Twilio also leverages AWS Direct Connect to establish private, dedicated network links between AWS and our 7 private Network Exchanges across the globe. Twilio Network Exchange (TNX) connectivity extends the reach of the AWS backbone for both our carrier interconnects and our customers on premise infrastructure to improve performance and security.



Twilio cloud infrastructure (built on AWS)

Built on a software-based infrastructure and distributed system, Twilio's platform handles concurrency, fast failover, and outage proofing that traditionally make hosting voice and messaging so complicated. With geographically distributed hosting, horizontally scalable architecture, and thousands of carrier connections worldwide, Twilio's cloud communications platform allows businesses to scale without implementing their own hardware.

Twilio's platform handles massive volumes of communications elastically and is provisioned to not exceed 60% capacity to sustain unusually high peaks of traffic and/or absorb failures. We are also cloud natives and are currently powering 140,000+ businesses and 100B+ interactions with a 99.95% SLA and 99.99% API availability.

Twilio's multi-tenant architecture leverages Edge Locations that are all designed to automatically scale, heal, and operate.

Twilio's platform handles concurrency, fast failover, and outage proofing that traditionally make hosting voice and messaging so complicated.



Super Network

Super Network leverages Edge Locations to directly connect with carriers, route media most efficiently, and provide redundancy and failover options. Twilio's Interconnect platform allows us to connect to our carrier partners through redundant fiber connections at our Edge Locations. This limits latency, jitter, shortens time to troubleshoot issues, and allows us to consume many more products and services, such as dual-use phone numbers from providers, that wouldn't be possible without a direct connection.

At scale, individual carriers have issues, frequently. The realities of the telecommunications industry—degraded service and/or outages are commonplace when looking at a group of carriers. Twilio makes use of multiple carriers for each route and operates a 24/7 Network Operations Center that is constantly monitoring the carrier networks and Twilio's dedicated communications engineers can optimize for changing traffic patterns.

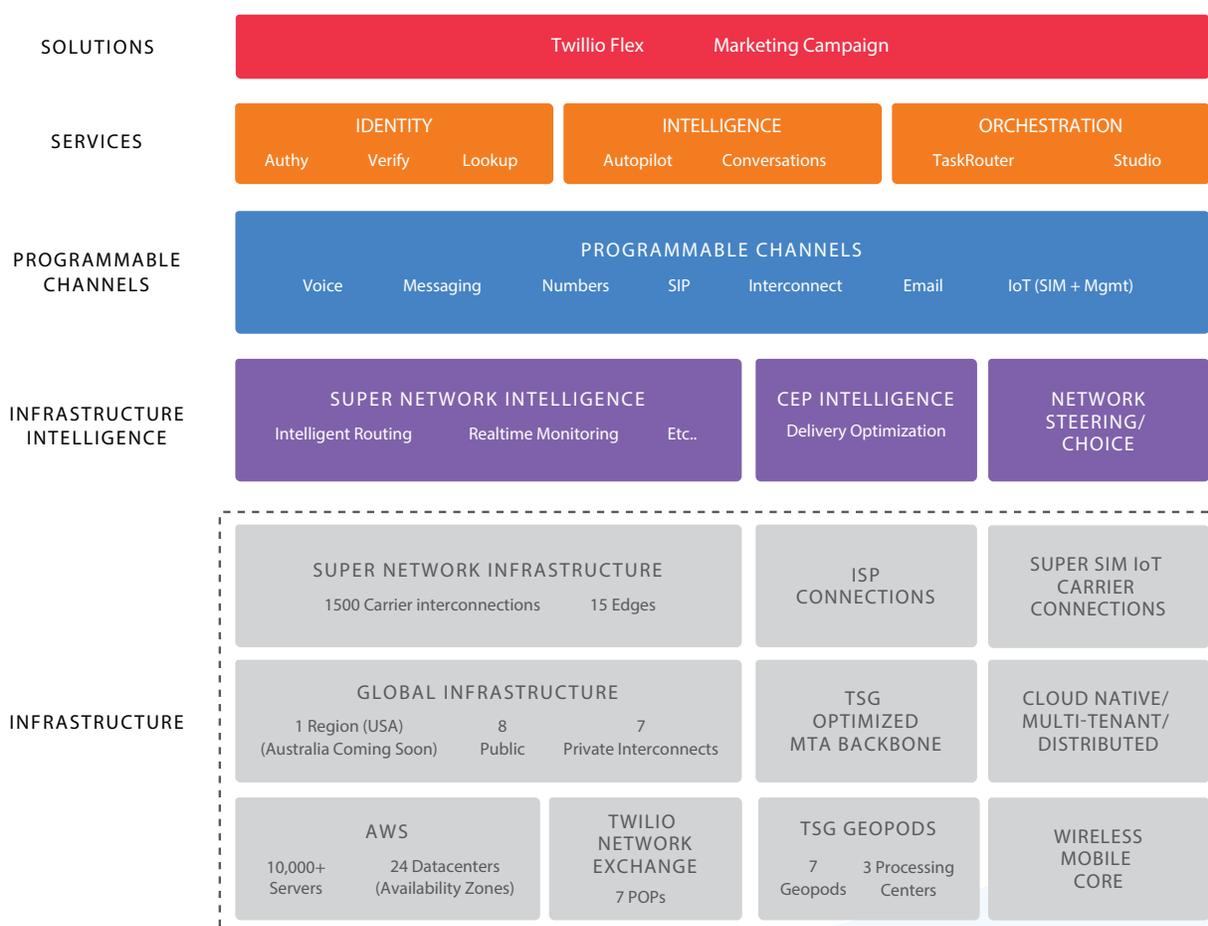
With massive volumes of traffic we are able to detect issues, often before our customers or carrier partners. We receive real time deliverability feedback from customers operating one time password use cases and our Feedback API, as well as our own internal Verify product. This traffic gives us real time feedback on handset deliverability over a number of carriers and destinations. We use this data to make our own routing decisions and provide status updates to our customers to assist them in making application decisions. This data-centric approach to our telecom architecture is what separates us from traditional PSTN carriers, and allows Twilio to deliver a best in class solution to its customers.

Twilio also offers detailed quality monitoring and unrestricted visibility into networking and connectivity issues across Twilio and third-party systems available to customers through Voice Insights. Voice Insights track more than 30+ dimensions of call quality in real time including caller ID preservation, post dial delay, silent call legs, day over day changes, changes, packet loss, and jitter. We continually monitor and test our network, and are able to catch 98% of incidents before customers and automatically reroute traffic.



Physical hardware infrastructure (built by Twilio aka Interconnect)

Twilio also manages and maintains dedicated hardware connections to AWS, our carrier partners, and our customers via our Interconnect Platform. Available now in the majority of our Edge Locations, we maintain highly available, dedicated private hardware to create a hybrid network that has improved reliability and resilience for customers. Our team of network engineering specialists are able to offer support and expertise for customers looking for more control and customization over their connection to Twilio.





Access to Twilio infrastructure

Twilio's Interconnect platform supports point to point redundant cross connections that physically terminate on different network devices. The redundancy is achieved at network layer 3 by running BGP over the point to point links. Twilio Interconnect is supported in all of our Edge Locations today ex Brazil. In each of these regions, media is kept locally and only signalling relays are sent back to US-East 1.

Interconnect also allows our customers who need better QoS, reduced latency and jitter, or who have heightened security requirements due to their use case or industry to connect directly to Twilio. Some customers choose to use Twilio Interconnect rather than send unencrypted traffic over the public internet in order to secure their communications, have more visibility into trouble shooting, increase throughput, and ensure a more consistent network experience.

Twilio offers a variety of private connectivity options for customers to align with their existing architecture design and application requirements:

Cross-connect: Cross-connect offers a dedicated physical interconnection with Twilio's Edge locations for customers that want a private internet connection for security and quality reasons. Each physical cross-connect that links the customer to Twilio has redundant fiber links to separate availability zones in AWS.

Bring your own MPLS: Bring your own MPLS is available to customers who want to avoid the public internet for security and/or quality reasons. Customers need an existing MPLS carrier relationship to send traffic through this private connection. Twilio supports a cross-connect between a customer's MPLS and one of our Edge locations.

VPN: A VPN allows for an encrypted virtual point to point connection over the internet for customers who have heightened security concerns but do not need the lower latency and quality that comes from private connectivity. Customers can establish point to point connection with Twilio via IPSec tunnels over the public internet to ensure that media and/or signalling cannot be intercepted.

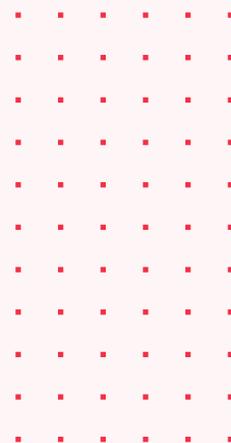


Software resiliency & testing standards

Competitors talk about new product/feature release every 6 week, Twilio releases new products/features/enhancements every 2 days. How are we able to achieve this velocity while maintaining our quality SLAs?

Twilio's platform is built using a 6-dimension engineering framework based on industry standards for product maturity that we call the Operational Maturity Model (OMM). This defines the standards for our engineering and ensures that our products are "well-hardened" for growth and scale. Ironman status is the highest ranking for each dimension.

Competitors talk about new product/feature release every 6 weeks, Twilio releases new products/features/enhancements every 2 days.





Security

The first dimension of OMM is security. All engineers at Twilio have passed secure code training, have proper authentication and authorization. Cloud communications are operated with the latest security practices used by the largest banks and multinational companies. Strong encryption, role-based access control, and signature validation protect the experience. Dedicated security experts are always improving security policies and best practices.

Architecture

The second dimension is architecture. Led by our team of internal architects, we develop detailed sequence diagrams for all flows through our services, including subsystems and their dependencies.

All of our products and solutions are built with GDPR and business continuity in mind. We offer data handling strategies for sensitive data including message body redaction, observe data retention periods, and allow customers to encrypt and delete recordings as well as manage sign on and access to their applications.

Data encryption at rest and in transit: Encryption is managed at the product level, but, generally, customer data stored in Twilio's AWS infrastructure will be encrypted at rest using industry-standard algorithms (AES-256). There are currently no products for which Twilio will provide blanket encryption-at-rest guarantees, beyond what is offered to customers at the product level.

- **Messaging:** Customers can control phone number and message body data once messages have been sent or received if [Message Redaction](#) is enabled.
- **Voice:** [Voice Recording Encryption](#) is a feature that allows customers to encrypt their recordings with a public key. Once the feature is activated, only the customer will be able to decrypt the recordings. There is no one at Twilio that will be able to decrypt the recordings

Data in transit within Twilio's private VPC network is not universally encrypted. However, we do encrypt data at the Edge when it is sent outside the VPC (usually with TLS). [A VPC is a private network inside of Amazon made up of a private block of IP addresses that you can only route traffic to through very specific mechanisms.] Twilio supports TLS 1.2 to encrypt network traffic between the customer application



and Twilio. Customers need to enable TLS on their side. Twilio will default to the highest cipher supported by the customer's client. Twilio is not able to guarantee that all connections to our carriers partners are encrypted.

Customers are responsible for managing any personal data they process via the Twilio Services. Twilio provides customers with a number of self-service features, including the ability to delete, retrieve, or restrict use of their content. Once deleted, the data may be retained on backup systems for up to 30 days — sometimes longer, depending on the service the customer is using. Our API documentation goes into more detail; for example, if you're using the SMS service, you can view our docs [here](#) for the minimum time-to-live of each element of an SMS message. For Voice, you can review our docs [here](#). After termination of the agreement, any content that customers do not delete will remain in place for 30 days, after which we will automatically initiate deletion, unless we are required by law to retain it for a longer period.

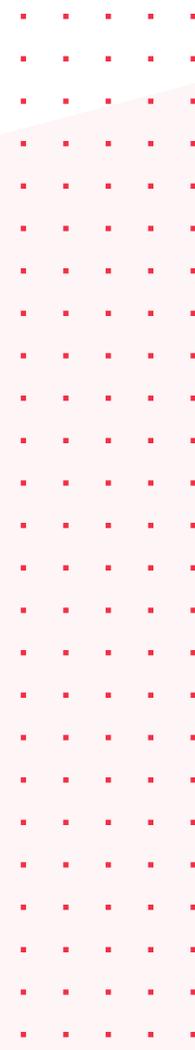
Twilio complies with GDPR & Twilio Programmable Voice is PCI DSS Level 1 certified.

Message redaction is applied to message bodies and numbers.

For HIPAA, we require volume level encryption for all PHI data at rest, no matter how long the data is at rest, but we don't require field-level encryption until you hit 7 days unless it is also considered sensitive PII.

Resiliency

Resiliency is critical and we run load tests to determine baselines, test scaling to 10x expected capacity to determine bottlenecks and instance ratios, and identify failure modes to determine safe operating parameters. We also leverage proper retry logic to ensure logarithmic back off and overall limit on retries.





Build/Test

Once our code has been developed, we run thorough testing and deployment management procedures to ensure seamless transitions. We follow proper source code and version management best practices, and our system makes use of automated workflow management tools to enable deployment testing and automated rollback. We have also deployed automated testing across separate development, staging, and production environments.

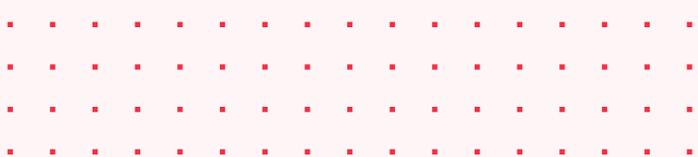
We test for failure isolation to guarantee that the failure of one dependency doesn't affect another, and have designed our systems to ensure that there are no single points of failure, and use longevity testing under load to detect leaks, resource starvation and degradation in performance over time.

Operations

We have defined SLA's for all of our services. Our runbooks are up to date and enumerate all operations. To track and manage these SLAs, we have instrumentation in place that includes system dashboards and alerting is setup. We also have a defined on-call schedule and pager rotations. Capacity planning has been completed for all of our services and performance throughput has been documented in our runbooks.

Support

We create product documentation and tutorials for customers in our Knowledge Base and documentation which cover errors and guidance on how to troubleshoot. For further questions, customers can reach out to our Customer Support team who has been trained on all of our products. We also have dedicated incident response teams proactively monitoring and patching potential vulnerabilities around the clock. Lastly, we keep customers up to date on the latest status of all our products on the Twilio Status Page.





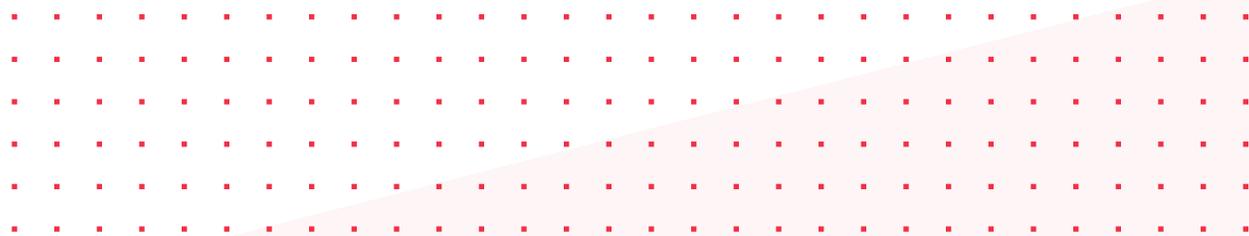
Product development lifecycle: Twilio develops products in a 4 phase lifecycle. Alpha products are for internal use and testing only, and the teams become familiar with the OMM principles they will need to fulfill.

Private Beta products have limited availability to primarily friendly customers. They do not have an SLA and have undergone only an initial OMM evaluation. The team incorporates learnings and feedback throughout the private beta.

Public Beta products are broadly available and allow customers to opt-in. The products are also undergoing OMM re-evaluation based on the implementation of learnings from the Private Beta phase.

We continually evolve our offerings to increase the capabilities, footprint, and product availability we offer globally.

Once a product has reached Ironman status on all OMM dimensions, it moves into GA. GA products are characterized by having a public SLA.



CONCLUSION

Twilio's Global Infrastructure provides the foundation for the Customer Engagement Platform. By providing a global communications infrastructure designed around resilience, redundancy, and low latency, developers can build any engagement solution, across all major channels, without compromising deliverability, or worrying about managing complex carrier relations and traditional telecommunication networks. By measuring each aspect of the Global infrastructure across an Operational Maturity Model (OMM), Twilio ensures that products and services are well hardened for growth and scale, and deliver 99.999% API uptime.

Twilio's Global Infrastructure is used by more than 172,000+ customers - from digital disrupters, to the world's largest enterprises, services phone numbers in 100+ countries. With continued investment, Twilio is delivering more edge locations to improve performance, and offer a variety of options for private connectivity including cross-connect, bring your own MPLS, and VPN to ensure that the needs of the world's most demanding customers are met.

If you'd like to learn more about the Twilio Customer Engagement Platform, and Global Infrastructure, we encourage you to talk to an expert [now](#). We can't wait to see what you build.



Questions?

Talk to your Twilio sales contact.
We can't wait to see what you build!

[Contact sales](#)



© 2020 Twilio