Everything You Need to Know About Gremlin Security

> *Gremlin provides an extremely simple to use API for setting up automated fault injection testing on your services.*
>
> **Michael Wong, Twilio**

We break things. That's our job.

To do this, you install our code directly onto your systems—yes, even your production systems—to deliberately try and undermine them.

Naturally, you have security questions. But rather than just saying, "trust us," we've put together a Q&A that covers all the things we do to ensure your data stays secure. In this document, we've tried to anticipate and answer your questions, but if you have any additional concerns, please don't hesitate to contact us.

Summary

The Gremlin architecture is designed from the ground up to be secure. We start with the architecture. Gremlin has three primary components: a client (often called an "agent"); the service layer, or API; and the web-based user interface. We protect each of these components with unique security mechanisms depending on its particular risks, including encrypting all communications, utilizing MFA, and strict change control procedures. We back up these claims through regular external security auditing and penetration testing.

When it comes to customers' data, we adhere strictly to principles of least privilege. We only access the data that is required to run our attacks. We never access or collect sensitive or proprietary information.

For all these reasons and more, customers like Twilio, Expedia, Confluent, and Remind trust Gremlin to deploy chaos engineering principles to test the robustness of their systems. But if you have any questions about security, just ask. Our internal security experts will answer promptly and thoroughly.

Frequently Asked Questions about Gremlin Security

Q: What is the Gremlin architecture?

A: Gremlin has three primary components: a client (often called an “agent”); the service layer; and the web-based user interface. Each of these components has its own security risks, and therefore, unique protections.

Q: What exactly is the Gremlin client? What does it do?

A: The Gremlin client is a program that is inserted directly into your host systems. It is designed to be a tool for YOU to test the robustness of your system; to provoke such things as network failures, process failures, or poor latency due to unusual resource usage.

Q: What kind of access to my network and systems does the Gremlin client component require to do its job?

A: We believe strongly in the principle of least privilege. Because of that, our client only needs very fine-grained permissions to the resources it requires. No root permissions are ever needed.

Q: Can you give an example of the types of things the Gremlin client will require permission to do?

A: Our client will need to have the permission to impact the network by adding latency, or the permission to control process management so it can kill processes as a way of testing the robustness of your system. For this, we use Linux capabilities such as “cap_sys_time,” “cap_sys_boot,” “cap_net_admin,” “cap_kill.” We do not need access to any application layer data, and the Linux capabilities that Gremlin employs are documented on the security pages of our website. We also support application layer failure injection, in which case you would have the ability to define exactly what data was provided to the Gremlin client.

Q: Does the Gremlin client require any inbound network connections?

A: No. You never need to open additional ingress ports on your network. All network traffic is outbound, using transport layer security (TLS) on port 443 with strong 256-bit advanced encryption standard (AES) as following the NIST (National Institute of Standards and Technology) recommendations for data protection.

Q: Can the client run arbitrary commands on the host system?

A: No. The client does not have the ability to run arbitrary or ad hoc commands on the host system. All attacks are pre-planned, and hardcoded in the client, and restricted to only the actions we introduce based on strict parameters. In other words, the client cannot attack anything except what you intend it to attack.

Q: What data does the Gremlin client collect from the system?

A: The client can only collect information necessary to impart the requested attack, ensure the desired impact is achieved, and facilitate targeting. We do not function in the application layer or collect data from the CPU or memory. We therefore don't have access to the data that the system is processing. Targeting data is collected and surfaced back to the user as infrastructure available to attacks on, and can consist of: the IP address of the host, the container ID's of any docker containers the client has access to, and if run within AWS, the region and availability zone the client resides in (to facilitate targeting all resources in a particular region or AZ).

Q: What happens if my environment is not completely secure?

A: The Gremlin threat model assumes that your environment is properly secure. The Gremlin client cannot protect your host system from inherently poor security configurations. However, we have put additional safety "guardrails" in place to prevent malicious manipulation of the environment from affecting the client—for example, path rewriting—if a lack of security allows a malicious actor to change path information.

Q: What is the Gremlin service layer?

A: The service is the second tier of the Gremlin architecture. It is a REST API that acts as the control plane for attacks. This API delivers attack requests to the client from the user interface and receives details from the client about the attack performed, which are then reported back to the user interface.

Q: How do you secure the API?

A: As with the Gremlin client, all communication with the service API is via TLS connections providing strong (256-bit AES) encryption, and all data received is encrypted at rest using the same technologies. We provide a strong access system, ensuring that the only the right people have permissions to launch or stop attacks. We offer multiple robust login and authentication methods, including passwords with multi-factor authentication (MFA); security assertion markup language, or SAML (a standard for exchanging authentication and authorization data between security domains); and Google logins, with the ability to require MFA or SAML company-wide.

Q: What is the web app?

A: This third element in the Gremlin architecture is how you interact with Gremlin. It allows you to schedule attacks, perform administration tasks, and includes an important safety feature: a “kill” button that halts all attacks immediately and returns your system to normal state.

Q: How do you secure the Web App?

A: We realized the importance of ensuring that this interface is not compromised in any way. Because of this, we apply numerous security mechanisms to it.

First, as with all other aspects of Gremlin, all communication is via TLS connections providing strong (256-bit AES) encryption.

Second, as in the API and service level, we focus on authentication so that only the right people have permissions to launch or stop attacks. We offer multiple robust login and authentication methods, including passwords with multi-factor authentication (MFA); security assertion markup language, or SAML (a standard for exchanging authentication and authorization data between security domains); and Google logins, with the ability to require MFA or SAML company-wide.

For added security, we also use content security policy headers (CSP), which prevent website-based attacks like XSS (cross-site scripting) and CSRF (cross-site request forgery). We also use certificate pinning, which allows the admin of a server to “pin” a certificate authority’s (CA) public key signature to a certificate, which is verified by the client (delivered via SSL extension). Additionally, time-based tokens in MFA add extra safety.

Q: What kind of data does Gremlin collect from my systems?

A: As we said before, we adhere strictly to principles of least privilege. We only access the data that is required to run our attacks. This includes network interface names and addresses if we are testing the network's vulnerability, and process or container names if we will be attempting to kill processes or attack containers. We also use email addresses as usernames. The Gremlin client collects some targeting data which is surfaced back to the user as infrastructure available to perform attacks on, and can consist of: the IP address of the host the client is running on, the container ID's of any docker containers the client has access to, and if run within AWS, the region and availability zone the client resides in (to facilitate targeting all resources in a particular region or AZ). We never access or collect sensitive information such as sensitive personally identifiable information (PII), medical information as protected by HIPAA, or credit card data as protected by PCI.

Q: How do you protect the data you collect?

A: All data is encrypted both in transit and at rest, using strong encryption (256-bit AES or stronger). Because of least privilege, only Gremlin employees with a proven business need to access data are granted access. We require MFA and VPN use with address filtering for access to all systems holding customer data. We also perform address filtering with a hard-coded address filter in our VPN for redundancy. Finally, we employ multiple redundant controls to ensure data isolation between customers, including ownership tagging.

Q: How do we know your own security mechanisms are robust enough to store our data?

A: Gremlin partners with an external security firm, Bishop Fox, to annually audit our security infrastructure and policies. This auditing includes extensive penetration testing. All vulnerabilities that are found are promptly fixed and re-tested by our auditors to ensure the vulnerability is mitigated. We can provide you with the signed results of the most recent Bishop Fox testing. Just ask.

We have thoroughly documented our corporate security policy and practices and rigorously enforce employee compliance. All new employees receive personal security training from the security staff, and their training is regularly refreshed. We use an industry-standard key management system to protect all keys. We pay special attention to training our employees on the risks of phishing and social engineering, as that is a weak link in most organizations' security practices.

Q: How do you assess security risks in your organization?

A: Gremlin utilizes qualitative risk management following the guidance of NIST SP 800-30. We assess our security risks continually—every time a change happens, and as we plan for future changes. We also do this as a routine part of our internal audits every quarter.

Q: What change-management practices do you follow?

A: Gremlin adheres to a strict development and infrastructure change-management policy that includes review by members of our security team for all changes to either production code or infrastructure. We also automate the detection of available patches and promptly update any vulnerable software.

Q: Do you employ an intrusion detection system?

A: Gremlin employs real-time intrusion detection systems that are automated at both the network and host (cloud) layers. All activity performed through the Gremlin system is logged, and we regularly audit and monitor for suspicious behaviors.

Q: What about wireless access at Gremlin?

A: Our customers tend to have two concerns about wireless access: How do we authenticate users, and what communications channel is used. All wireless access is secured by WPA2 using individually identifiable authentication identities. We can tie every login directly to each user, and we do not permit shared accounts. We do not store sensitive data on systems connected to wifi networks.

Q: What would you do if your systems get compromised?

A: We go the extra mile and are extremely transparent about all security incidents to our customers.

We have a strict policy to inform our customers of any verified security breach within 24 hours, and of any suspected security breach within 48 hours.

We have a documented procedure to manage security incidents if and when they occur. This procedure includes preserving all evidence, performing appropriate forensics, and reporting the incident to all affected customers. Additionally, in the event of a security incident we will always work with our customers to investigate and determine the impact and resolve any outstanding issues related to the event.

Q: What guarantee do we have that you won't pull data out of our systems?

A: This point is covered in our contracts with customers, but additionally, our independent third-party auditors annually review our source code to ensure it does not perform data exfiltration.

Q: What do you do if one of your attacks gets out of control?

A: We—and, more importantly, our customers—have the ability to kill any attack, and halt any attack immediately, and revert to a known "good state." In the case where the client loses communication with the API it will automatically cease any active attacks and revert those impacts, this is referred to as a `dead man switch`. We understand that in the wrong hands, our tools could be used to cause a lot of damage. That is why we focus on the entire system when protecting our customers. Not only does the Gremlin client have those safeguards, but the API and user interface have multiple safeguards to prevent unauthorized access. (Those all impact the client, since the API is what controls the clients, and the UI is the primary way customers interact with the API.)

We understand that any failure that causes a business disruption to our customers equates to them losing money. In some cases, very significant money. Because of that, we are extremely careful to not unintentionally negatively impact the host system.