

Security Information Package

Online Voting Security for the 2018 Ontario Municipal Elections



Security Overview

Simply Voting Inc. will be providing the Internet and Telephone Voting System for the 2018 Ontario Municipal Elections. The Simply Voting system is secure and protects the secrecy of your vote.

Secret Ballot

Whether you use the internet or telephone to vote, your vote is instantly encrypted and stored with no possibility of your vote being traced back to your identity, just like a traditional paper ballot. It is impossible for municipal staff, Simply Voting employees or any other person to see how you have voted. Election officials will only be able to see that you cast your vote, the time you voted and the IP address or telephone number you voted from.

One Person One Vote

Only registered voters on the municipal list of electors will be authorized to access a ballot. Once you vote, using either internet or telephone, you are “crossed off” the list and cannot vote again. Even if you switch between the internet and telephone, even if you try to vote using several devices at the same time, the system will only accept one single ballot from each voter.

Auditing

Municipalities will assign independent auditors of the Internet and Telephone Voting System. Simply Voting provides designated auditors with appropriate access to observe that the system allows voting under proper circumstances and prohibits voting under improper circumstances. The Auditor may continuously monitor voting activity before, during and after the voting period. It is also impossible for the Auditor to see how you have voted.

For some municipalities: Once your vote is cast a receipt code is issued. Only you will know this code. Print this code or write it down. After voting has ended, you can look up your receipt code to verify that your vote was counted and (for some municipalities) how your vote was counted.

Protection Against Computer Hackers

Simply Voting is an expert in internet security and goes to great lengths to protect the voting system. All communications between your computer and the voting website are encrypted to ensure confidentiality. The internet ballot is tamperproof and there are multiple layers of security to protect the servers against attacks.

Protection Against Imposters

To vote, you will need to enter a password. Every password is a nine-digit numeric PIN that will be mailed to each voter in a Voter Information Letter prior to the "Voting Period". These PINs are randomly generated by Simply Voting and are printed, machine folded in security-tinted envelopes, and mailed directly to voters using Canada Post. As an added security measure, voters will also be required to enter their date of birth (or year of birth for some municipalities) to complete the voting procedure. Therefore, if your Voter Information Letter ends up in the wrong hands, another person will not be able to cast your vote without your PIN and your date of birth.



Technical Information

Top-Notch Security

Simply Voting was designed from the ground-up to eliminate the risk of electoral fraud or breach of secrecy:

- ⤴ Voters who bypass authentication or have already voted are denied access to the ballot.
- ⤴ One-vote-per-voter is guaranteed by marking electors as voted and storing the vote in a single transaction. Even if a voter submits the ballot simultaneously on several devices, this technology guarantees that only one vote is accepted.
- ⤴ Ballots are rigorously checked for validity before being accepted.
- ⤴ All administrator and voter activity is logged with timestamp and IP address in an Activity Log.
- ⤴ Communication between the voter's computer and our website is encrypted with TLS 1.2 and strong cipher suites to protect against current and future encryption attacks.
- ⤴ Our servers are "hardened" and are subjected to daily Trust Guard PCI Compliance security scans.
- ⤴ Our voting system is regularly subjected to penetration tests by Spirent SecurityLabs and source code security audits by HP Fortify.
- ⤴ Simply Voting adheres to guidelines established by the Open Web Application Security Project.
- ⤴ Any change to the voting system must pass an internal security review before going live.
- ⤴ All staff workstations are kept up-to-date and protected by access password, firewall, anti-virus, anti-spamware and disk encryption.
- ⤴ We use Domain Keys Identified Mail and the Sender Policy Framework to protect voters from phishing attacks.
- ⤴ Our servers are protected by a very powerful firewall, FortiGate Unified Threat Management, which includes an Intrusion Detection System and a redundant firewall on hot standby.
- ⤴ Simply Voting uses *CloudFlare* to protect against Denial of Service (DoS) attacks. *CloudFlare* has the most sophisticated mitigation technology on the market and has successfully blocked the largest DoS attacks seen on the internet. CloudFlare will be "always on" for the voting website. More info about their solution is available at <https://www.cloudflare.com/ddos/>
- ⤴ We use redundant *Anycast DNS* deployments which protects against DNS-based DoS attacks.

Fully Hosted & Reliable

Simply Voting is built on an enterprise-class cloud computing service powered by high performance IBM hardware, with full redundancy across the entire infrastructure (no single points of failure). Our data centre is in a stable mountain zone, away from earthquake, hurricane, tornado, and severe weather zones. The data center contains advanced power, cooling and security infrastructure, and Cisco Data Center 3.0 network architecture. It is staffed 24x7, backed-up by an offsite network operations center. We also use several Anycast DNS clusters to ensure fault tolerance at the DNS level.

Simply Voting uses third party offsite monitoring tools to automatically monitor key “vital signs” of our voting system 24x7 and a technical staff member is immediately notified of any anomaly. Simply Voting maintains a Disaster Recovery Plan as well as a Hot Site at a backup data center in a different geographical area. The Hot Site is synchronized with the primary data center using remote database replication. Should the primary data center experience an outage, we have the capability of quickly redirecting traffic of the entire voting system to the Hot Site, minimizing disruption to ongoing elections and avoiding any loss of data. You can rest assured that your election is always protected and available in the case of a disaster.

For telephone voting, Simply Voting uses industry leader Plum Voice as a voice-to-web interface layered on top of our online voting system. Every component in the Plum Voice, fault tolerant infrastructure has a backup and Plum's platforms have been tested by billions of calls since 2000. Plum's PCI Level 1 compliant operation actively secures and protects applications and data from digital, physical, and social intrusion vectors. There is no artificial cap on “ports”, the telephone voting system can handle spikes of millions of simultaneous calls at once.

Skyhigh Enterprise-Ready Rating



Simply Voting received the highest CloudTrust Rating from Skyhigh Networks. Skyhigh performs objective and thorough evaluations of cloud services based on a detailed set of criteria developed in conjunction with the Cloud Security Alliance (CSA). Services designated as Skyhigh Enterprise-Ready fully satisfy the most stringent requirements for data protection, identity verification, service security, business practices, and legal protection.

Confidentiality

Simply Voting takes secrecy of the vote very seriously. Votes are stored without any information that could be traced to an elector, so it is logically impossible for election organizers or even the Simply Voting system administrators to determine what a particular voter has voted.

All voter information is removed from our servers if you choose to have the election deleted. We never make use of voter information for anything other than voting and never share such information with third parties. Our privacy policy (available on the Simply Voting website) and voting system have been independently certified by TRUSTe for compliance with their Privacy Certification and Trusted Cloud requirements.



SSAE 16 Certification



Simply Voting is SOC 2 certified. The SOC 2 is a widely-recognized auditing standard issued by the American Institute of Certified Public Accountants (AICPA). An auditor's report details a service provider's ability to offer adequate controls and safeguards when they host or process data belonging to their customers. The audit focuses heavily in the areas of compliance, security and access. It addresses topics such as backup and recovery, computer operations, and human resources.

Our primary data center, TeraGo, is similarly SOC 2 Type II certified and our secondary data centre, iWeb is also SSAE 16 SOC 1 certified.

These certifications are an independent validation of the quality, integrity and reliability of Simply Voting's infrastructure and services.



Advanced DDoS Protection Service

With CloudFlare

Denial-of-service (DoS) attacks are on the rise and have evolved into complex and overwhelming security challenges for organizations large and small. Although DoS attacks are not a recent phenomenon, the methods and resources available to conduct and mask such attacks have dramatically evolved to include distributed (DDoS) and, more recently, distributed reflector (DRDoS) attacks—attacks that simply cannot be addressed by traditional on premise solutions.

CloudFlare's advanced DDoS protection, provisioned as a service at the network edge, matches the sophistication and scale of such threats, and can be used to mitigate DDoS attacks of all forms and sizes including those that target the UDP and ICMP protocols, as well as SYN/ACK, DNS amplification and Layer 7 attacks. This document explains the anatomy of each attack method and how the CloudFlare network is designed to protect your web presence from such threats.

Below you will find detailed information on these attacks and how the CloudFlare network protects against them:

- Layer 3/4 attacks
- DNS amplification attacks
- SMURF attacks
- ACK attacks
- Layer 7 attacks
- Making DoS a thing of the past

Layer 3/4 attacks

Most DDoS attacks target the transport and network layers of a communications system. These layers are represented as layers 3 and 4 of the OSI model. The so called “transport” layer of the network stack specifies the protocol (e.g., TCP or UDP) by which two hosts on a network communicate with one another. Attacks directed at layers 3 and 4 are designed to flood a network interface with attack traffic in order to overwhelm its resources and deny it the ability to respond to legitimate traffic. More specifically, attacks of this nature aim to saturate the capacity of a network switch, or overwhelm a server's network card or its CPU's ability to handle attack traffic.

Layer 3 and 4 attacks are difficult—if not impossible—to mitigate with an on-premise solution. If an attacker can send more traffic than a network link can handle, no amount of additional hardware resources will help to mitigate such an attack. For example, if you have a router with a 10Gbps port and an attacker sends you 11Gbps of attack traffic, no amount of intelligent software or hardware will allow you to stop the attack if the network link is completely saturated.

Very large layer 3/4 attacks nearly always originate from a number of sources. These many sources each send attack traffic to a single Internet location creating a tidal wave that overwhelms a target's resources. In this sense, the attack is distributed. The sources of attack traffic can be a group of individuals working together, a botnet of compromised PCs, a botnet of compromised servers, misconfigured DNS resolvers or even home Internet routers with weak passwords.

Because an attacker launching a layer 3/4 attack doesn't care about receiving a response to the requests they send, the packets that make up the attack do not have to be accurate or correctly formatted. Attackers will regularly spoof all information in the attack packets, including the source IP, making it look as if the attack is coming from a virtually infinite number of sources. As packet data can be fully randomized, even techniques such as upstream IP filtering become virtually useless.

With CloudFlare, all attack traffic that would otherwise directly hit your server infrastructure is automatically routed to CloudFlare's global Anycast network of datacenters. Once attack traffic is shifted, we are able to leverage the significant global capacity of our network, as well as racks-upon-racks of server infrastructure, to absorb the floods of attack traffic at our network edge. This means that CloudFlare is able to prevent even a single packet of attack traffic from a traditional layer 3/4 attack from ever reaching a site protected by CloudFlare.

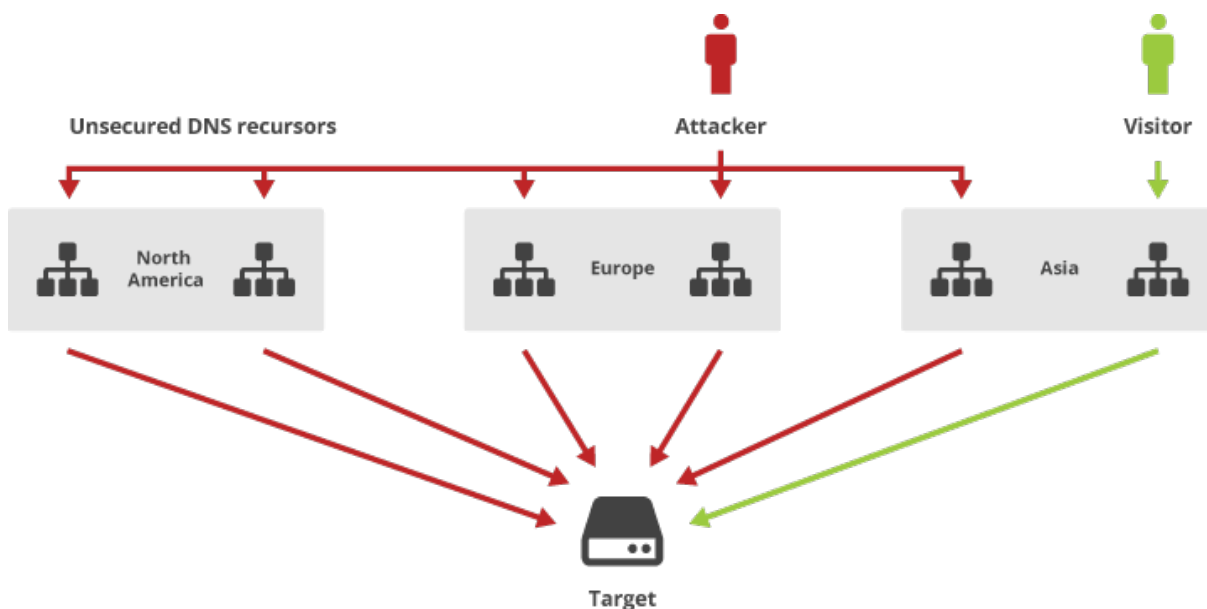
DNS amplification attacks

DNS amplification attacks, one form of DRDoS, are on the rise and have become the largest source of Layer 3/4 DDoS attacks. CloudFlare routinely mitigates attacks that exceed 100Gpbs, and recently protected a customer from an attack that exceeded 300Gpbs—an attack the New York Times deemed the “largest publicly announced DDoS attack in the history of the Internet.”

In a DNS reflection attack the attacker sends a request for a large DNS zone file—with the source IP address spoofed as the IP address of the intended victim—to a large number of open DNS resolvers. The resolvers then respond to the request, sending the large DNS zone answer to the IP address of the intended victim. The attackers' requests themselves are only a fraction of the size of the responses, allowing the attacker to amplify their attack to many times the size of the bandwidth resources they themselves control.

DNS Reflection Attack Without CloudFlare

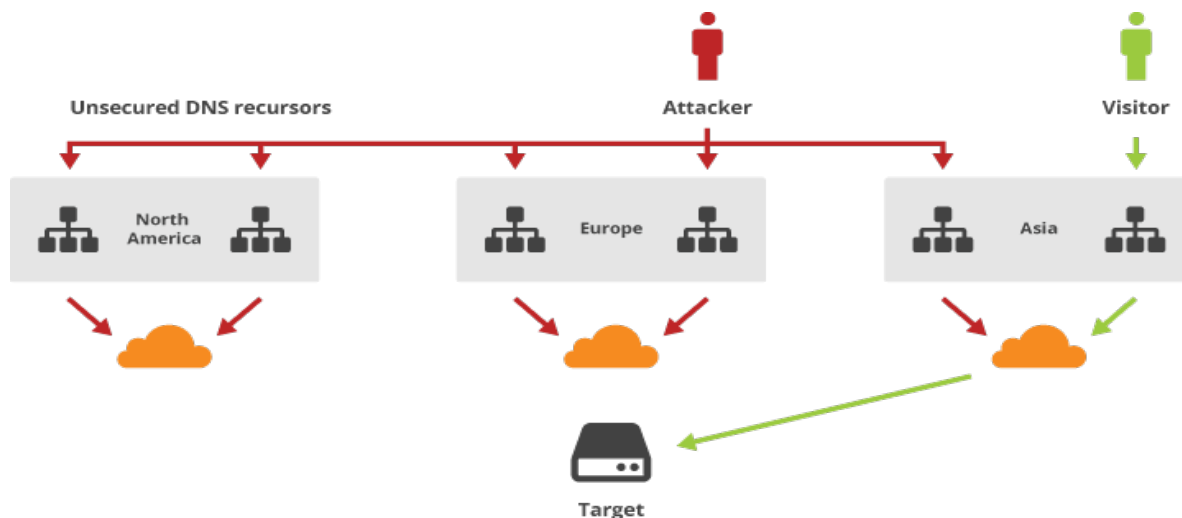
An attacker gathers resources, like botnets or unsecured DNS recursors, and imitates the target's IP address. The resources then send a flood of replies to the target, knocking it offline.



DNS Reflection Attack with CloudFlare

An attacker gathers resources, like botnets or unsecured DNS recursors, and imitates the target's IP address. The resources then send a flood of replies to the target, but they are blocked regionally by CloudFlare's data centers. Legitimate traffic can still access the web property.

There are two criteria for an amplification attack: 1.) a query can be sent with a spoofed



source address (e.g., via a protocol like ICMP or UDP that does not require a handshake); and 2.) the response to the query is significantly larger than the query itself. DNS is a core, ubiquitous Internet platform that meets these criteria, and therefore has become the largest source of amplification attacks.

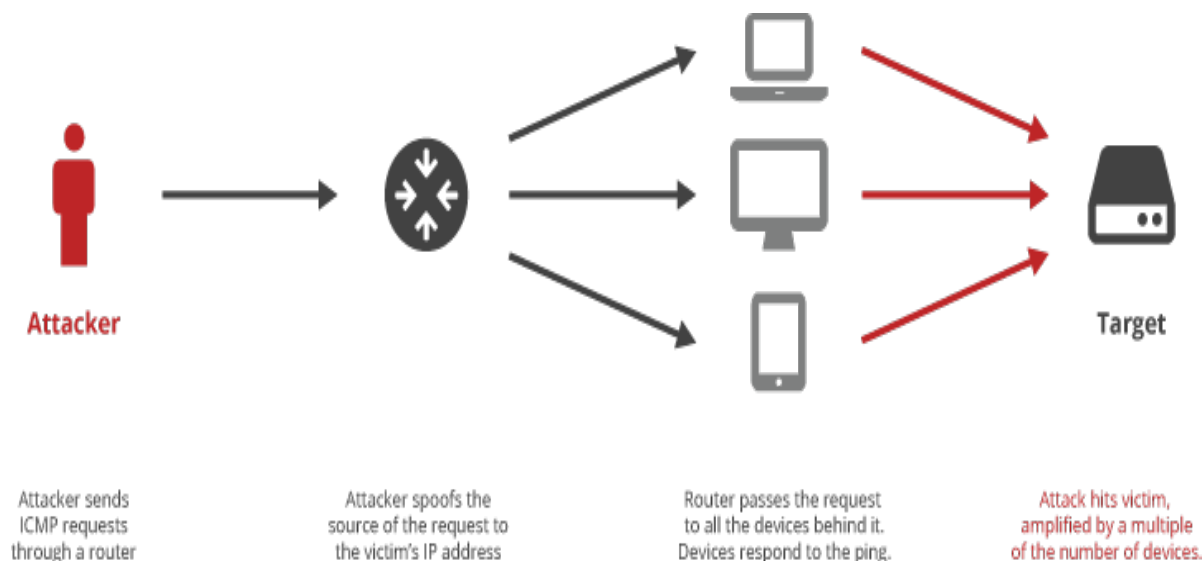
DNS queries are typically transmitted over UDP, meaning that, like ICMP queries used in a SMURF attack (described below), they are fire-and-forget. As a result, the source attribute of a DNS query can be spoofed and the receiver has no way of determining its veracity before responding. DNS is also capable of generating a much larger response than query. For example, you can send a (tiny) query (where x.x.x.x is the IP of an open DNS resolver): That's a 64-byte query that can result in a 3,223-byte response. In other words, an attacker is able to achieve a 50x amplification over whatever traffic they can initiate to an open DNS resolver.

CloudFlare's "Anycast" network was specifically designed to stop massive layer 3/4 attacks. By using Anycast, we are able to announce the same IP addresses from each of our 23 worldwide data centers. The network itself load balances requests to the nearest facility. Under normal circumstances this helps us ensure that your site's visitors are automatically routed to the nearest data center on our network to ensure the best performance. When there is an attack, Anycast serves to effectively scatter and dilute attack traffic across our entire network of data centers. Because every data center announces the same IP address for any CloudFlare customer, traffic cannot be directed to any one location. Instead of the attack being many-to-one, it becomes many-to-many with no single point on the network a single point of failure.

SMURF attacks

One of the first amplification attacks was known as a SMURF attack. In a SMURF attack an attacker sends ICMP requests (i.e., ping requests) to a network's broadcast address (i.e., X.X.X.255) announced from a router configured to relay ICMP to all devices behind the router. The attacker then spoofs the source of the ICMP request to be the IP address of the intended victim. Because ICMP does not include a handshake, the destination has no means of verifying if the source IP is legitimate. The router receives the request and passes it on to all the devices that sit behind it. Each of these devices then respond back to the ping. The attacker is able to amplify the attack by a multiple equal to the number of devices behind the router (i.e., if you have 5 devices behind the router then the attacker is able to amplify the attack 5x, see the diagram below).

SMURF attacks are largely a thing of the past. For the most part, network operators have configured their routers to disable the relay of ICMP requests sent to a network's broadcast address.



ACK attacks

In order to understand an ACK attack, one must delve into the world of TCP. When a TCP connection is established there is a handshake. The server initiating the TCP session first sends a SYN (for synchronize) request to the receiving server. The receiving server responds with an ACK (for acknowledge). After that handshake, data can be exchanged.

In an ACK reflection attack, the attacker sends lots of SYN packets to servers with a spoofed source IP address pointing to the intended victim. The servers then respond to the victim's IP with an ACK creating the attack.

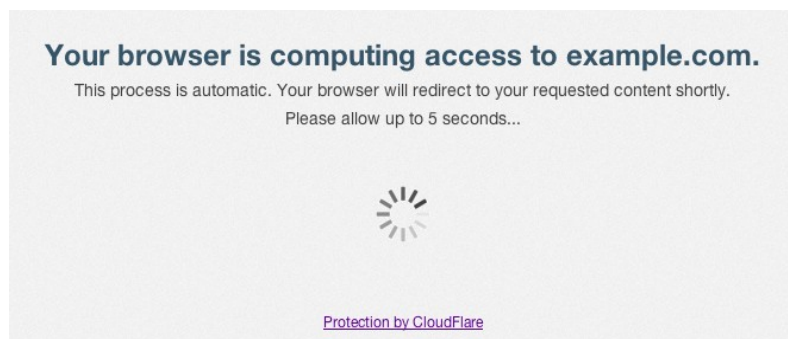
Like DNS reflection attacks, ACK attacks disguise the source of the attack making it appear to come from legitimate servers. However, unlike a DNS reflection attack, there is no amplification factor: the bandwidth from the ACKs is symmetrical to the bandwidth the attacker has to generate the SYNs. The CloudFlare network is configured to drop unmatched ACKs, which mitigates these types of attacks.

Layer 7 attacks

A new breed of attacks target Layer 7 of the OSI model, the “application” layer. These attacks focus on specific characteristics of web applications that create bottlenecks. For example, the so-called Slow Read attack sends packets slowly across multiple connections. Because Apache opens a new thread for each connection, and since connections are maintained as long as there is traffic being sent, an attacker can overwhelm a web server by exhausting its thread pool relatively quickly.

CloudFlare has protections in place against many of these attacks, and in real world experiences we generally reduce HTTP attack traffic by 90%. For most attacks, and for most of our customers, this is enough to keep them online. However, the 10% of traffic that does get through traditional protections can still be overwhelming to customers with limited resources or in the face of very large attacks. In this case, CloudFlare offers a security setting called “I’m Under Attack” mode (IUAM).

IUAM is a security level you can set for your site when you’re under attack. When IUAM is turned on, CloudFlare will add an additional layer of protections to stop malicious HTTP traffic from being passed to your server.



While a number of additional checks are performed in the background, an interstitial page is presented to your site’s visitors for 5 seconds while the checks are completed. Think of it as a challenge where the tests are automatic and visitors never need to fill in a CAPTCHA.

After verified as legitimate by the automated tests, visitors are able to browse your site unencumbered. JavaScript and cookies are required for the tests, and to record the fact that the tests were correctly passed. The page which your visitors see when in IUAM can be fully customized to reflect your branding. I’m Under Attack mode does not block search engine crawlers or your existing CloudFlare whitelist.

Making DoS a thing of the past

As technology advances, DoS attacks will only increase in complexity and magnitude. Traditional on premise DoS solutions simply cannot adapt to the wide range of new attack vectors, and are rendered completely ineffective for attacks that exceed an organization’s network capacity.

The CloudFlare network is designed to mitigate and keep pace with the changing threat landscape. CloudFlare, as an operator of one of the largest global networks on the Internet, is able to leverage its aggregate network capacity across 86 points of presence, and is able to learn from attacks against any individual customer to protect all customers on our network.



The Data Center

Security Advantages

Physical Security


The TeraGo Kelowna 'GigaCenter' data center employs multiple physical security measures – including mantraps, proximity pass and biometric. There are seven layers of security between the front door and an individual computer rack. Physical security measures include:

- Single story concrete building
- Security Cameras strategically located throughout the facility
- Multiple pan-tilt-zoom cameras outside the facility
- Camera images are recorded, searchable and archived for a minimum of 90 days
- Proximity pass and biometric scanners, supporting multiple security zones
- Motion sensors and intrusion detection sensors
- Intrusion sensors with audible alarm
- Steel doors and two-stage man traps
- Computer racks are individually locked
- Access control system is secured in a room accessible by authorized personnel only
- Manned and monitored security desk
- Security systems are monitored 7x24 by both the on-site NOC and an off-site third party

Security Processes

TeraGo employs extensive security processes that support our clients' needs. Our processes have been audited by a third party by evidence of our SAS70 certification. Existing, audited processes include:

- All entrances are locked at all times; Two factor authentication (badge and biometric) is required for access to the facility and to the data halls
- Background and criminal record checks for all employees and authorized (badged) contractors
- All employees must wear a photo-ID badge at all times while in the facility
- Each employee and authorized contractor must badge in when arriving, and badge out when leaving the facility (no tailgating); a perpetual log is maintained of what personnel are onsite
- Badge and biometric access is controlled in zones, ensuring personal have access to authorized areas only
- Changes to access are documented and approved by management
- The ability to create, modify or delete access authorization is restricted by management
- Processes are in place to remove access when an employee or contractor is terminated or a badge is lost
- The access control system is logged, searchable and archived; logs are retained for at least 90 days
- Visitors are required to sign a visitor log, provide a government issued photo ID, and wear a visitor badge while in the facility
- Visitors are escorted at all times
- TeraGo personnel are on-site 7x24x365



The TeraGo Chief Security Officer is responsible for the ongoing review, management, optimizing and documenting the Security Plan.

Systems & Data Security

All TeraGo staff have RCMP criminal records checks before attaining facility access accreditation and access to TeraGo systems. Systems and Data security processes are confidential and proprietary. These processes have been reviewed by an independent third party auditor, as evidenced by TeraGo's SAS70 certification.

Network Resiliency & Security

The GigaCenter is carrier-neutral and is served by numerous major Canadian telcos, including Shaw, Bell, Rogers, Allstream and Telus. TeraGo also manages a private 10Gbps network into major Canadian cities, including Vancouver, Calgary and Toronto, and Seattle Washington. This provides access to high capacity, low cost bandwidth from major North American and International communications providers. Data on the TeraGo private network is transported at layer 2. TeraGo can cross connect with the client's preferred telco at the peering points mentioned above, providing a secure and highly available service.

Clients can procure network services directly from their preferred telco, or purchase network services from TeraGo and our telco partners. Clients can use their own VPN and encryption technologies to support their security requirements, for services delivered from the GigaCenter. Fiber feeds into the GigaCenter are delivered through diverse underground conduits into the facility.

The data center LAN is fully redundant with 10Gbps capability to every cabinet and device within the facility. Built on the Cisco Nexus Data Center 3.0 platform, it has no single points of failure and supports concurrent maintenance. All customer data traffic is isolated on private VLANs within our GigaCenter switching with separate layer 3 routing interfaces created per VLAN at the core routing layer. No layer 2 traffic is carried between VLANs, and logical layer 2 VLAN segments are never shared between clients.

A portfolio of optional and highly recommended services is available to enhance our clients' security programs:

- Firewall & Network Security Gateway
- Includes IPS and IDS
- Includes customer portal for self management and customization
- Web Security service
- Web Application Security service
- Email Security service



Security of Cloud Services

As described above, a key element of our Cloud security is the separation of client's data traffic at layer 2, on both the TeraGo private WAN and GigaCenter LAN. Each client has their own private VLAN, and VLAN segments are not shared between clients.

Other key elements include:

- Each Cloud VM has its own RAM
- Each Cloud VM has its own SAN space, allocated privately in a secure multi-tenant model
- Virtualization is delivered using industry leading VMware vSphere

Cloud server infrastructure is delivered from a TeraGo owned and managed high security vault within the GigaCenter facility

For more information, please write to info@simplyvoting.com