# Technology Trends

## Zero Trust Networking

Enterprise Architecture, Chief Technology Officer Branch

Version 0.1
Date 2019-5-8

Shared Services Canada | Services partagés Canada

Canada

# Table of Contents

# Business Brief

*Zero Trust Networking* (ZTN) refers to a data networking architecture model first named by John Kindervag of Forrester in 2010, many parts of which have been in use for some years. The total cost of cybercrime (direct impact, but also mitigation) is expected to ramp up from USD3 trillion in 2015 to USD5 trillion by 2020. Most current approaches to data-, information-, and network security revolve around *perimeter defense*, also known as the *castle and moat* model. In this model, the outer perimeter of the organization is defended (as in the moat of the castle), with data allowed in/out through a very limited set of points (the *drawbridge* of the castle) that are implemented as inbound and outbound firewalls. The assumption (which used to be valid) is that threat actors are on the outside, while those on the inside of the perimeter can be *trusted*.

The present-day threat reality is considerably different for a couple of reasons:

- Most organizations currently encourage a significant amount of *mobile computing* as well as *bring your own devices* (BYOD), while increasing amounts of data and applications residing in the cloud. These have completely eroded the notion of a perimeter.
- *Threat actors* have become more sophisticated: in addition to the traditional external hacker, various attacks now gain access to witting or unwitting insiders to assist in a security breach. After initially breaking into one system, attackers are usually able to *move laterally* to various other systems. The last years has seen the rise of *advanced persistent threats* (APTs) that encompass all of this threat tradecraft.

These factors imply that threat actors are *already inside* the organization. A first attempt at mitigating the issue is a *dual trust boundary* in which servers containing the organization's "crown jewels" are protected with an additional layer – though distributed data and the cloud make this insufficient nowadays.

The ZTN model is currently the best available approach for defense[1], and revolves around a number of principles:

- All *resources are accessed in an authenticated manner*. All such accesses are authenticated using some notion of *identity*. This applies to humans, apps or other processes accessing data, other processes and services, or computing resources. This mitigates possible insiders or lateral movements by threat actors.

[1] It is worth noting that Forrester no longer "owns" this approach in the sense that numerous vendors provide zero trust solutions. Furthermore, competing analysts Gartner have created a much more elaborate model known as *Continuous Adaptive Risk & Trust Assessment* (CARTA), which partly includes the features of zero trust.

- All *resources are accessed securely*. When access is authorized, the information exchanged is securely encrypted. This mitigates eavesdropping by threat actors already inside the organization.
- *Least-privilege* (needed to get the job done) is given to any identity needing to access resources. This limits the impact of any stolen identity.
- *Log and inspect* all network traffic, server, and identity activity. Coupled with *analytics*, this ensures that an audit trail exists for detecting unusual attempts to access resources.

All of these principles are self-evident, and some may be in use in organizations, though a comprehensive ZTN approach requires integrating their use, as well as the governance and change management aspects.

# Technical Brief

Thankfully, ZTN relies heavily on *already existing technologies*, meaning that the technical risks and costs are well understood. Those technologies and architectural approaches are:

- *Identity and Access Management* (IAM). Adequate identity management is obviously crucial to ensuring that all access to resources is authenticated/authorized. Such resources may be data, but also computing services, devices, etc. Information systems are already designed around the notion of identity for human users, but IAM for ZTN requires that all possible users of a resource also be identified and authorized. This implies IAM for processes, for example one piece of software making use of another, or *Internet-of-Things* (IoT) devices feeding into a *data lake*. Such use cases will enormously increase the number of identities being managed – well into the billions when IoT is included. Traditionally, such non-human identification (and sometimes human identification) has been done on the basis of identity proxies such as IP address – something that is clearly vulnerable to attack as they can be spoofed. Modern IAM uses certificates, signatures, and multifactor authentication for robustness; only recently has computing power become cheap enough to support these on a large scale. These techniques are all in-place at SSC, though primarily for employees, and they potentially need upgrading over the coming years to mitigate threats from quantum computing.
- *Micro-segmentation* of networks. Segmentation of networks is a well-known technique for separating sub-networks – originally for performance reasons and then for security reasons. When security became a motivation, the separation between two network segments was effectively done with a firewall that inspects network traffic between segments, logging problematic traffic, and perhaps allowing only traffic from certain IP addresses, or conforming to specific protocols. With ZTN, the need arises to do this at fine granularity – essentially each machine (or at most a small number of like-functionality machines) is in its own micro-segment. Furthermore, the inter-segment firewall *must* restrict traffic to the minimal set of protocols (traffic types) required. If the machine providing a service does not provide authentication, then this internal firewall also does the required IAM.

In no case is an IP address sufficient for authentication. These internal firewalls require a significant ability to log and report more centrally (see below). Predictably, placing each machine in a micro-segment will lead to a proliferation of internal firewalls, though computing power has reached the point that they can be implemented in software as opposed to rack-mounted hardware in a data-center. In the case of virtualized or cloud based infrastructure, the firewalling and micro-segmenting are part of the *Software Defined Network* (SDN), further lowering the cost of implementation. SDN-type implementations have one significant caveat: they require highly secure virtualization *hypervisors*. A breach of an insecure hypervisor would cripple ZTN and give access to all data and traffic on at least that specific physical machine.

- *Ubiquitous encryption*. Since threat actors are effectively inside the "perimeter" of the organization and may snoop on network traffic, all interactions across the network should be encrypted. Additionally, data-at-rest should be encrypted to mitigate break-ins to a particular database or file server. Modern algorithms, CPUs and accelerators have made such large-scale encryption trivially fast.
- *Dynamic and conditional policies*. The principle of least privilege is a cornerstone of ZTN. Over the employment of a particular person or the active period of a process, their required privileges/access will vary depending on task, and such privileges should be kept at the most restricted level that still enables the work[2]. Current practice at most organizations is to raise privilege level as needed (perhaps even setting it "high" initially) and almost never lower it. ZTN changes that to being much more dynamic, and lowering it again once a specific task has been completed. When applied to processes or devices (not humans), this can be implemented via micro-service architectures/APIs – the fact they support only a single service means they can execute with least privileges.
- *Logging, inspection and analytics*. Awareness of threat actors is still needed, even in ZTN. This is best enabled by extensive traffic inspection (using the micro-segments) and logging. Rather than inspecting logs manually, they are typically processed by data-science/analytics (for security) and machine learning algorithms to show anomalies that indicate attack attempts. The resulting analytics are used to generate reports and visualizations as required for governance of ZTN.

# Industry Use

There are few figures on industry uptake of ZTN, though it is universally agreed to be the best current strategy for *information security*. The top-tier tech savvy banks, social media (LinkedIn, FaceBook), online shopping (Amazon), and infrastructure (Apple, Google,

[2] This interacts strongly with change management and human resources issues, where employees inevitably feel frustrated or stymied in executing routine work efficiently. For that reason, some latitude is needed in the interpretation of "least privilege."

Microsoft) all implement some amount of ZTN, though sometimes only in areas of their most confidential customer information or financial systems.

ZTN products are also available from the leading network infrastructure, cloud services, and encryption providers.

# Canadian Government Use

As mentioned in the *Business Brief*, most components of ZTN are in use somewhere already. In particular, defense, intelligence and policing (DND, CSE, CSIS, and RCMP) will already have extensive ZTN in-place, also at the interpersonal (not just IT system) level, using the principle of need-to-know and least privilege.

# Implications for Shared Services Canada (SSC)

## Value Proposition

The SSC value proposition is set against the backdrop of the new realities of:

- More sophisticated threats: insiders but also APTs, some of which may be nation-state actors
- No perimeters at the Government of Canada (and SSC and its clients in particular), thanks to: geographical distribution, encouragement of work-from-home, mobile devices, BYOD, SmartCities and IoT, and most impactfully, extensive cloud usage.

That is, threat actors are typically already inside the organization, thanks to attacks and lateral movement of insiders.

The most obvious value proposition to SSC and clients is lowering the cost incurred by cybercrime and cyberattacks – a cost that is a function of *risk* times *impact* (in dollars). The precise amount is difficult to quantify, but ZTN lowers both factors in the cost.

ZTN is easy to roll out *incrementally*: with each re-architecting of infrastructure (for example, while moving to the cloud), ZTN can be incorporated as a core implementation decision. The scalability of such an incremental approach is especially important with increasing use of IoT by SSC clients – which will eventually lead to billions of devices interconnecting.

Lastly, the cost of IT governance can be lowered via ZTN, as clear boundaries of access are set, and considerably more data is logged regarding access to resources.

## Challenges

Various components of ZTN are already in-place within SSC, and the first challenge are two aspects of change management: moving away from assuming a perimeter and

trusting all insiders (whether people or systems); making this move without damaging SSC-employee relationships in which trust is implicitly part of valuing employees. Roll out timing for ZTN would present related challenges of prioritizing some resources (including employees) for ZTN, while encouraging appreciation of the security win.

Since identity management is integral to ZTN, existing SSC concepts of identity will need adjustment. Currently, *Internal and External Credential Management* (ICM and ECM) extend the notions of credentials (user name and password) towards true identity management, and the most important challenges will come in allowing *devices to have identity* (e.g. IoT devices) and also scalability (towards billions of IoT devices).

On a technical level, SSC's extensive moves to cloud infrastructure can allow for easily implementing ZTN, but will depend on very high quality hypervisors for virtualization. The entirety of ZTN fails if a threat actor successfully breaks into a hypervisor. Encryption and authentication of most or all data and access to resources will also be computationally intensive and therefore budgeted into hardware requirements.

# Considerations

SSC already implements the ZTN components to varying degrees and consideration is needed for the roadmap to maximizing the impact of this. Change management deserves the most attention, as it will directly affect the mindsets of employees at SSC and at clients: thinking in terms of threat actors already being (virtually and digitally) inside the organization. As SSC increases that awareness, SSC should also consider how to retain employee commitment when zero trust at least superficially gives the message of "not trusting anyone" and requires the principle of least privilege. This is also a good opportunity to consider improving the quality and efficiency of reporting and governance, as enabled by ZTN.

Logistically, SSC's cloud commitment presents the ideal time and vehicle for integrating ZTN – all analysts agree that systems should be *ZTN by Construction* (ZbC) as opposed to added as an afterthought. SSC should therefore (re)consider the cloud roadmaps to include ZbC. Because ZTN requires additional encryption and authentication, this will have implications for costs of the cloud move.

Lastly, consideration must be given to extending the current identity management infrastructure (ICM and ECM) so that devices and processes become *identified actors* in the broad sense within SSC, and their permissions, access and activities are appropriately managed and logged.

# References


- "What is Zero Trust? A model for more effective security." CSO from IDG. https://www.csoonline.com/article/3247848/what-is-zero-trust-a-model-for-more-effective-security.html

- "Zero Trust is an Initial Step on the Roadmap to CARTA." Gartner, 10 December 2018. https://www.gartner.com/doc/reprints?id=1-641B4AK&ct=190114&st=sb
- "Defend Your Digital Business From Cyberattacks Using Forrester's Zero Trust Model." Forrester, 12 September 2018. https://reprints.forrester.com/#/assets/2/716/RES61555/reports
- "Zero Trust Networks." Doug Barth and Evan Gilman, July 2017. O'Reilly Publishing.