PaperCut

A PAPERCUT WHITEPAPER

Edge Mesh: How to get print resiliency and security without servers

October 2020

Introduction **3**

On-premise - printing locally **3**

Cloud print today **4**

Enter: Edge Mesh! **5**

What features come with “the best of both worlds”? **6**

Converting from local to cloud 6

Self healing components 7

Full-system resiliency 7

Maximum security 8

Principles of Edge Mesh’s zero trust 10

Looking to the future **11**

Works cited **12**

Thank you **13**

Authors 13

Significant contributors 13

Introduction

It should come as no surprise that businesses love the cloud - the benefits are clear and organizations are quickly shedding their locally hosted applications. Printing, on the other hand, is something that most businesses today still rely on "On-Premise" solutions for.

There's really something to be said about local infrastructure. It's great to be able to unplug the internet and still have things work for you. No matter what the application, there are still good arguments in favor of keeping things local, and good arguments for pushing up to the cloud. Despite some great features in the cloud, it isn't always a perfect option- a solution that truly offers the best features of both cloud infrastructure and local resources.

Printing is one of those things that historically has not made technical leaps outside your local network. There are ways to get some of the awesomeness of the cloud into your printing environment, but sometimes you might find yourself missing that sweet simplicity of the local alternative. In a printing environment, we'd really like something that truly provides the best of both cloud and local infrastructure. That something exists, and it's called **Edge Mesh.**

> 72% of businesses expected to increase their cloud adoption this year...
>
> **Internet Legal Technology Association**

On-premise - printing locally

Printing today is firmly rooted in the traditional client-server relationship within a local network. This long history of local printing means that managing a local environment is familiar. A local print environment isn't locked to specific brands or product lines to function. A physical server simply stays on-site to provide access control, and ease of deployment. All of your core functionality is centralized in your network, and it remains fully manageable at all times.

This is great! Local assets are quick to respond since network latency is practically eliminated. Meanwhile, hardware can be physically accounted for, and you can control everything directly. You know it's going to work because you can personally fix anything you see going wrong before it happens. The environment doesn't rely on a cloud provider to stay available. If things go bad in the cloud, you can rest easy, since your local resources will still be there, ready to meet your needs. Relying on local

network infrastructure is vital in making your printing resources stay as available as possible. Since people reasonably expect printing to just work, this stability makes for a great printing experience.

The practical benefits of control and availability aren't the only reason to rely on local printing. Keeping things local also has another vital advantage: data sovereignty. The data you send through your local infrastructure stays in your hands at all times. No other organizations have the opportunity to handle your information. But, of course, local printing comes with a cost. You need skilled IT people onsite or on call to manage and maintain local infrastructure.

Cloud print today

Printing in the cloud takes a lot of administration out of your hands. You don't have to worry about your server, because it's backed by dozens of servers all managed by the cloud provider (Fig. 1). The cloud offers a predictable budget with the service billing model, a hugely scalable environment, and a unique guarantee of availability. If an organization has a surge in printing, the cloud can automatically scale to meet their demands. All this is an excellent start to a great solution. However, you lose a lot of the control that comes with local infrastructure. For starters, an organization now needs a reliable Internet Service Provider. The organization's data sovereignty is reduced, and your environment now relies on a 3rd party to keep your printing running.

With cloud functionality as a foundation, many organizations look to achieve the benefits of local printing by keeping the job locally and only sending metadata to the cloud (Fig. 2). Solutions like PaperCut MF hosted in the Cloud, among others, add network-edge appliances to localize critical parts of the cloud's workload. This method is a great way to move toward the best of both cloud and local. No matter what you call it though, an edge appliance is still a piece of hardware for which your IT staff are the first line of defense. It's not quite a server, but it comes with some of the same maintenance and single point of failure concerns that a local server has.

Fig. 1

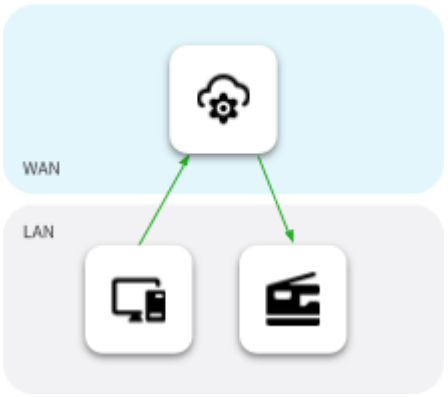

Fig. 2

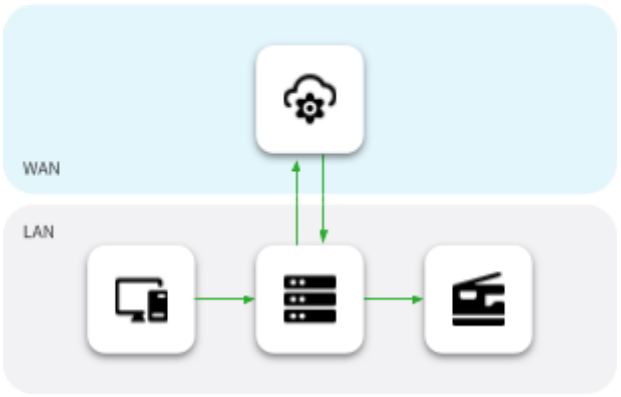

Our current options are either ease of use and affordability with cloud, full-featured and secure functionality with local infrastructure, or a little bit of both (but not quite everything we want!) with a hybrid environment.

Enter: Edge Mesh!

We use the philosophy of “Mesh” every day - even if we don’t think about it. Ask yourself: “Do we want to keep the knowledge of our operating procedures with just one person? Would we prefer a limited reliance on key personnel? What if an individual with valuable knowledge is out sick for the day when we need them?” The answers are easy: it’s better to have a team build in redundancy and avoid any single point of failure. The philosophy behind mesh is the same: enable the parts of your network to function as a team, each able to perform multiple roles.

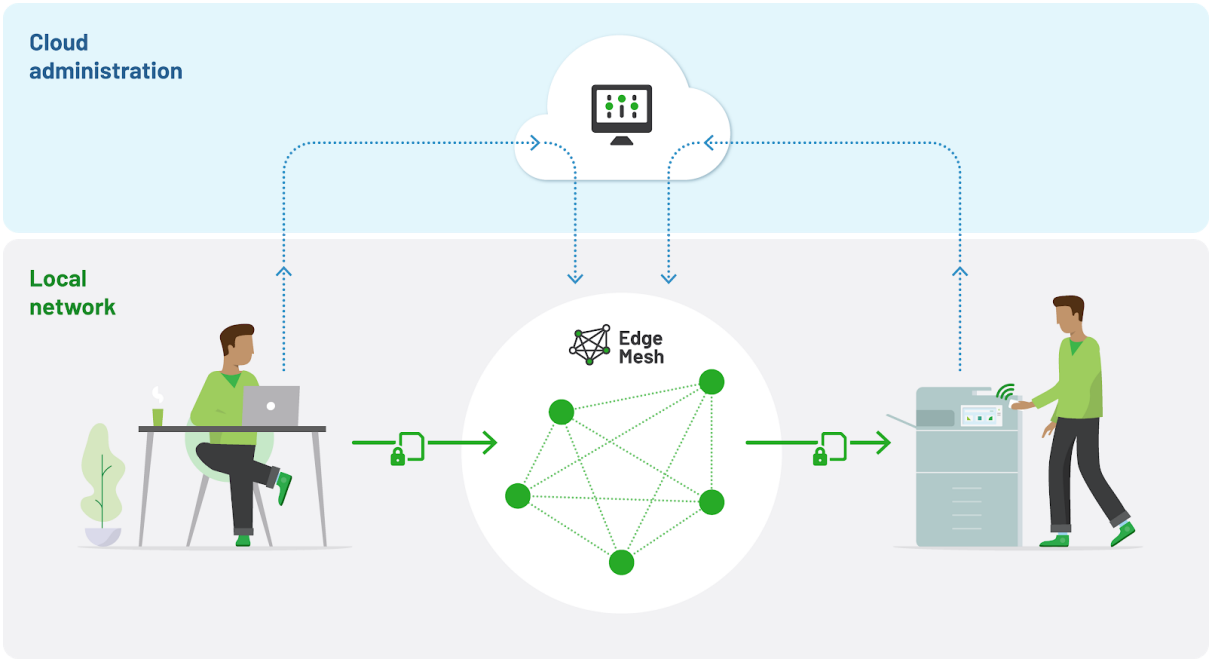

Edge Mesh solves the problems of both cloud printing and local infrastructure by turning your existing devices into redundant components of a collective print server.

To speak in more general terms, “Edge Mesh can be defined as a computing paradigm that uses (a) mesh network of Edge Devices and routers to enable distributed decision-making within the network.” The nodes within the mesh are the “brains” of the process. They learn how to interact with things from the cloud, but they are fully capable of making their own decisions once they know what to do. Each computer in the network becomes a “node” in the mesh - so rather than printing to a queue that you’ve set up on your server, the job is received by the software running in the Edge Mesh. The minute you press print, the Edge Mesh will work with the cloud to determine the best nodes for the job and distribute the file accordingly. Any node can go offline, but your printing system will stay available because of the inherent resiliency of the distributed, self-healing resources. This is the advanced interaction of local infrastructure and cloud computing that we’ve needed to truly get the best of both worlds.

What features come with “the best of both worlds”?

Designing Edge Mesh into a print solution allows organizations to have a true enterprise printing experience with the budget of a small to medium-sized business. High-availability, native zero-trust and, very high scalability are all built into Edge Mesh's architecture. With all of these local resources working together to replace the server as we know it - administrators can also rest easy without the worry of migrating and protecting their servers in the cloud.

Features	Local Infrastructure Printing	Public Cloud Printing	Edge Mesh
Offline Functionality	YES	✘	YES
Near-Zero Latency	YES	✘	YES
High Security	YES	✘	YES
Direct Management	YES	✘	YES
Scalability	✘	YES	YES
High Availability	✘	YES	YES
Cost-Effective	✘	YES	YES
Easy administration	✘	YES	YES
Off Premise Print Submission	✘	YES	YES

Converting from local to cloud

The migration process from physical infrastructure to cloud infrastructure is traditionally a taxing one. Building complex migration schedules, planning, and staging the move, shouldering the expense of a parallel migration - things can get messy. By contrast, the Edge Mesh eliminates this complexity. The mesh adopts your existing infrastructure and dynamically grows with your network. You can manage a partial or full migration on your own time. If an organization is not ready to say farewell to its server infrastructure, it can still utilize its server's computing power to benefit the Edge Mesh. Server nodes can be promoted to skew the node selection process toward them.

> For many companies, the cost savings alone can be a driver towards deploying an edge-computing architecture.
>
> **Network World, 2019**

Self healing components

In a network of 50, 100, or even 1000 computers, you might expect one or two servers to handle printing. However, with an Edge Mesh printing solution, this relationship is perfectly scalar. As a new host is added to the printing environment, it can commit some of its resources to the Edge Mesh; basically adding a server for every new host. As your printing environment grows, so does the number of hosts capable of handling the processed jobs. If a single node fails, another node in the Edge Mesh will take its place. As long as at least one node stays online, job processing continues.

The Edge Mesh's high availability doesn't end with hardware failover. The mesh also applies high availability to the jobs it processes. Jobs are triple-protected from the failure of any single node since, on submission, the document makes its way to three separate nodes. Maintaining job-redundancy allows admins to be sure that when a user releases a print job, the job will be there as efficiently as possible.

Full-system resiliency

Edge nodes open up many possibilities for resilient and fail-safe operation. For example, nodes within the Edge Mesh could communicate with one another to determine where the job is and where it needs to go. If a node doesn't have access to the printer that the job needs to be released on, the Edge Mesh can find a node that can send the job. With these mechanisms in place, it would take a very significant network disruption to stop printing (and, by then, you've probably got other things to worry about...).

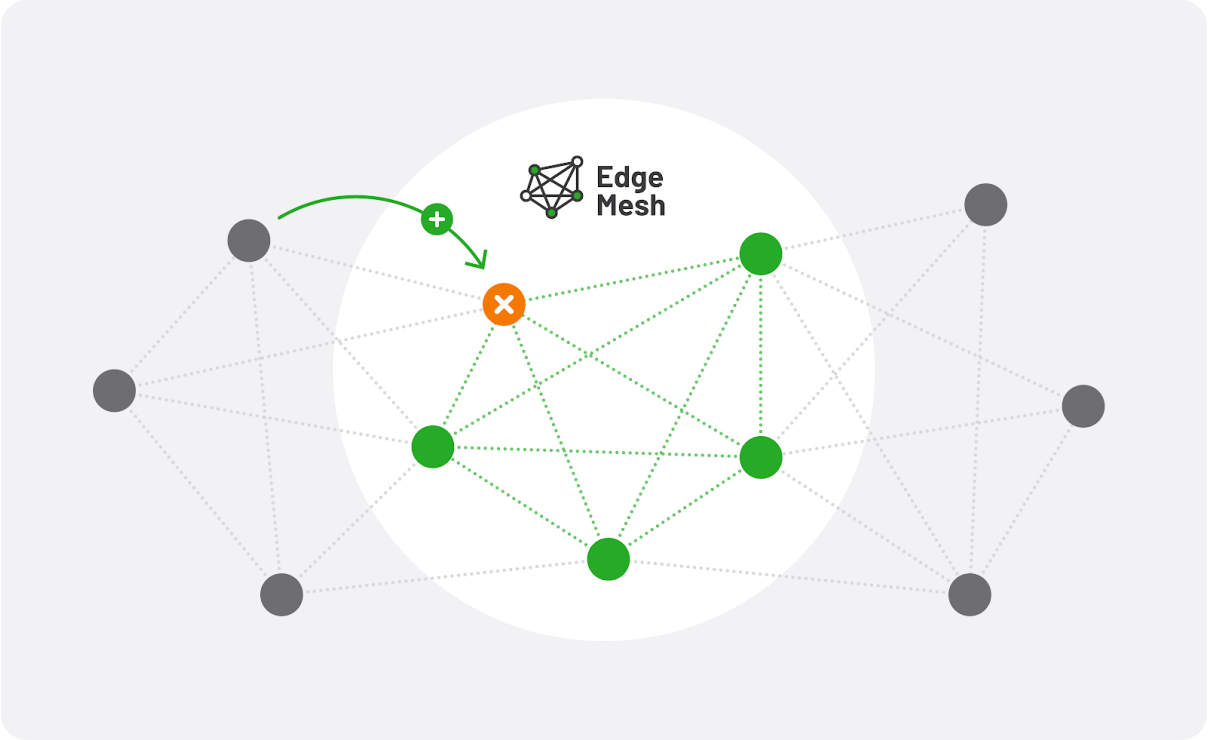

As your printing environment grows, so does the number of hosts capable of handling the processed jobs.

Additionally, network administrators don't have to worry about dropping availability to update if issues or inefficiencies are found within the software. In a traditional server environment, the downtime for the servers would need to be scheduled. The risk of losing print functions during the upgrade would just have to be accepted. With Edge Mesh, the nodes individually communicate with the cloud provider for updated information. This means that the nodes will patch themselves in a way that maintains enterprise availability for the entire print infrastructure. No downtime. No data loss.

Maximum security

Edge Mesh is also an important part of securing interactions with the cloud. Sending a job through the cloud opens up otherwise-secure documents to the risks of being "**on the internet**". You need to trust that your cloud provider is handling your data carefully and that all points that the job travels through are correctly maintained for security. Having a fully capable mesh of computing resources available locally means that a job never has to leave the network to begin with. Not to mention, in an Edge Mesh, no organization outside of your own ever handles your documents. Of course, just "being local" does not necessarily mean "being secure".

> "From a security perspective, Zero-Trust can better track and block external attackers, while limiting security breaches resulting from internal human error."
>
> **Defense Innovation Board, 2019**

"Average Size of a Data Breach: 25,000+ Records"

"Average Cost per-record: $242.00 USD"

IBM, 2020

Mesh networks are not just something that you can "plug into" like a physical network. The Edge Mesh doesn't rely on a physical connection or a single traditional password to enable access. Instead, access is governed by the cloud which challenges new hosts to verify access to the system. Since the cloud is the source of all the Edge Mesh's initial configurations, trust is built into the mesh without the potential of compromise.

An increasingly common practice with IoT device management, and printer infrastructure, is to host devices on a separate dedicated VLAN - often called an Infrastructure VLAN. This isolates devices into their own security zone (e.g. You may have a complex password on your WIFI, but connecting to the network using a printer's ethernet cable may make this security moot!).

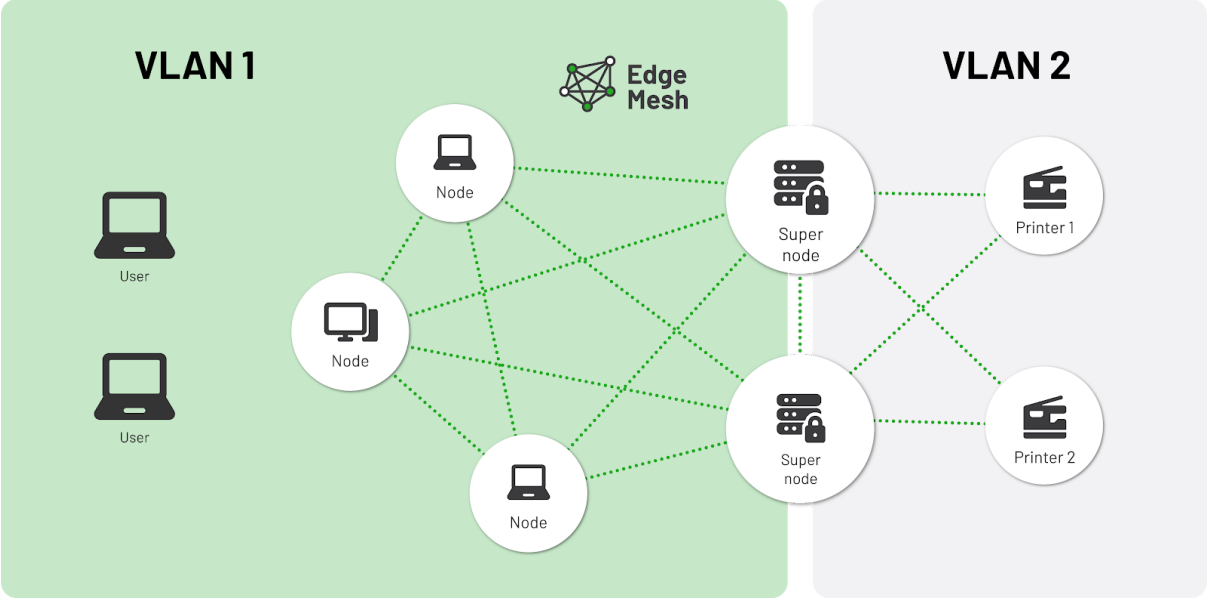

PaperCut's Edge Mesh helps out here too. The mesh contains an intelligent routing algorithm so jobs can "hop" via nodes in the Edge Mesh to navigate to other network segments.

By passing at least one node on the boundary between the VLAN, print jobs can flow. A common way to accomplish this is to set up a Virtual Machine running the Edge Node software that has access to both VLANS.

In a similar methodology used with Kerberos, the cloud will stand in as the manager of the relationships within the mesh. This means that the necessary information to decrypt the data in the mesh cannot be extracted from any single node on the network. Instead, the cloud can provide the foundation for the encryption.

From the ground up, our Edge Mesh is built to adhere to the principles of Zero-Trust networking. Even the metadata that does leave the network makes its way to the cloud using an encrypted HTTPS connection. Although, by definition, the principles of Zero Trust networking are not just an "authenticate once and call it a day" configuration. Once an Edge Node is added to an Edge Mesh network, it can't just go printing without extra validation and security measures. Ongoing communication between the Edge Mesh and the cloud enables the mesh to apply security settings to jobs in real-time. If a node is kicked out of the mesh, its token is removed, and the other nodes will reject authentication.

Principles of Edge Mesh's zero trust

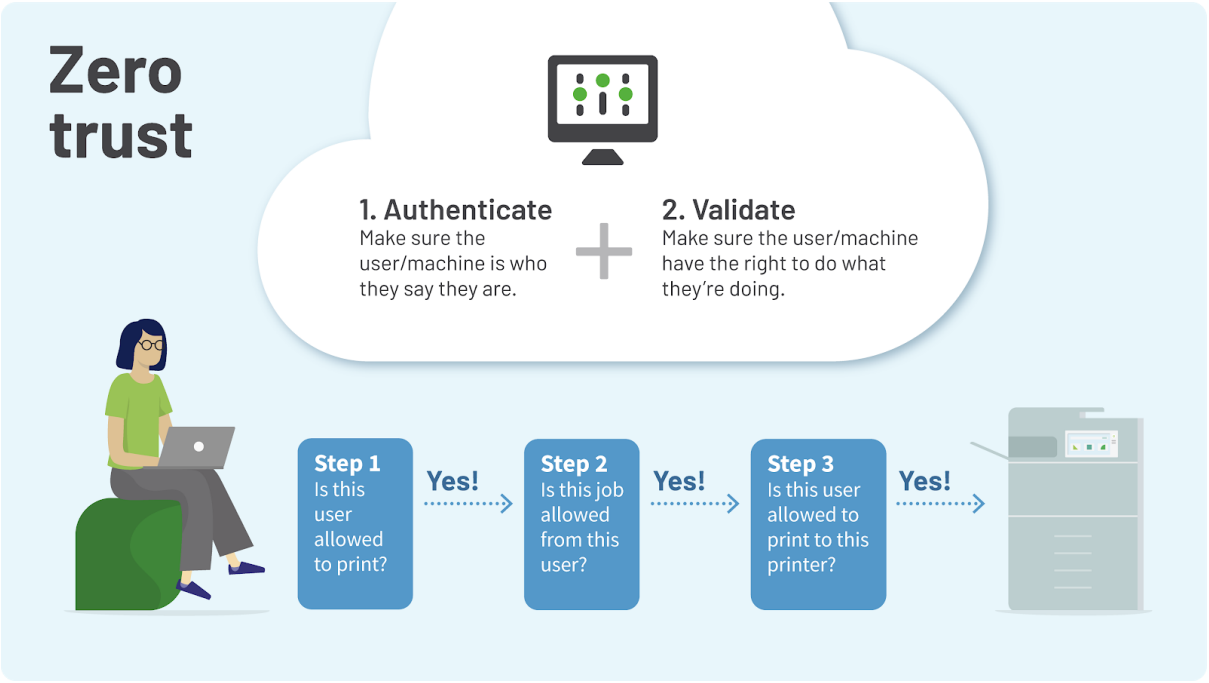

Every step in submitting a job is made with security in mind.

When a user sends a job, it's securely transmitted via HTTPS into the Edge Mesh. Each of the three copies of the job is securely distributed through the Edge Mesh to the best Edge Nodes for the job. After analysis, the Edge Mesh will encrypt the data using a three-part key provided by the cloud and store it for later. Unlike your car or a safe, the next thing the Edge Mesh does is discard the key. Even if an Edge

Node with a job stored on it has a sudden change of heart and becomes a bad actor, the data within may as well be gibberish. To unencrypt and process that job into meaningful information, the Edge Node needs to contact the cloud, authenticate, get the keys again, and unlock the file.

> 47% of enterprise IT security teams lack confidence in their ability to provide Zero Trust," despite 79% of organizations planning to adopt it.
>
> **Pulse Secure, 2020**

The Edge Mesh is built from scratch to comply with the most secure industry standards. Each Edge Node is purpose-built to offer a higher level of security than is commonly found in traditional print. Zero-Trust guards against all access-based compromises and protects data at every point within the Edge Mesh. Edge Mesh stands to make zero-trust accessible to any administrator- because it's built into the way the mesh works.

Looking to the future

Edge Mesh has the potential to shift how printers work in our network. Mesh technology stands to bring the practically-zero latency and high security of local network resources along with the massive scalability and enterprise high availability of the cloud. A company with just three or four computers in their entire network could never consider an enterprise, zero-trust, print environment - until Edge Mesh, that is. Despite the impressive feature set that can be offered with Edge Mesh technology, it doesn't come at the cost that's been traditionally associated with such performance. Edge Mesh stands as a democratizing shift in the structure of our networks. No longer is the barrier for entry of high availability insurmountable for the small-to-medium-sized business. For organizations of every size, Edge Mesh will drive performance and security and will be a significant first step in beginning to spread the use of high-security, zero-trust networking to all organizations. PaperCut believes that Edge Mesh is the future of managed printing services, and we're excited to be working with this game-changing innovation in our cloud platforms, PaperCut Pocket and PapertCut Hive.

Works cited

McCue, Jim "ILTA 2019Technologuy Survey Executive Summary" *International Legal Technology Association,* September 2019

"What is edge computing and why it matters" *Network World,* 2019

"The 2020 State of It" *Spiceworks* 2020 Spiceworks, Inc.

"The State of Application High Availability" *SIOS*

DelBene, Kurt, et al. "The Road to Zero Trust (Security)" *Defense Innovation Board,* 9 July, 2019

"2020 Zero Trust Progress Report" *CyberSecurity Insiders/Pulse Secure*, 2020

Y. Sahni, J. Cao, S. Zhang, and L. Yang, "Edge Mesh: A New Paradigm to Enable Distributed Intelligence in Internet of Things," in IEEE Access, vol. 5, pp. 16441-16458, 2017, doi: 10.1109/ACCESS.2017.2739804.

"Cost of a Data Breach Report 2020" IBM, 2020

Thank you

Authors

Jim Gebhardt (Systems Architect)

Tyler Alvord (Frontline Support Engineer)

Significant contributors

Geoff Smith (Distinguished Product Engineer)

Chris Dance (Founder & CEO)

Bez Mozaffari (Senior Product Engineer)

Irina Sutyagina (Product Engineer)

PaperCut HQ

sales@papercut.com
papercut.com

Support

papercut.com/support