

A PAPER CUT WHITEPAPER

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

June 2023





Introduction.....	3
On-premise - printing locally.....	3
Cloud printing today.....	4
Enter: Edge Mesh!.....	5
What features come with “the best of both worlds”?.....	6
Converting infrastructure from local to cloud.....	7
Self-healing components.....	7
Full-system resiliency.....	7
Maximum security.....	8
Secure local printing.....	9
Zero Trust networking.....	10
Principles of Edge Mesh zero trust.....	10
Looking to the future.....	11
Works cited.....	12
Thank you.....	13
Authors.....	13
Significant contributors.....	13
PaperCut HQ.....	13
Support.....	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 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, though, 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 the local network. Yes, 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. Printing environments really need 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...

International 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. To function, the local print environment isn’t locked to specific brands or product lines. 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 on site or on call to manage and maintain local infrastructure.

Cloud printing 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 the 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






So, our current options are 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. Well, the philosophy behind Edge Mesh is the same: enable the parts of your network to function as a team, each able to perform multiple roles.



The 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, the Edge Mesh is a computing paradigm that uses a mesh network of edge devices (users' computers) and routers to enable distributed decision-making for printing within the network.

The edge devices are called “edge nodes” within the mesh and are the “brains” of the print management process. They learn how to interact with the cloud, but they are fully capable of making their own decisions once they know what to do.

Each computer in the network becomes an “edge node” in the Edge Mesh - so rather than printing to a queue that you’ve set up on your server, you print to the software running on an edge node (for example, a user’s PC).



The second you press print, the Edge Mesh works with the cloud to determine the best nodes for the job and distributes the file accordingly to ensure it's successfully printed. If a node happens to go offline, your printing environment stays available because of the Edge Mesh. This reliability is the advanced interaction of local infrastructure and cloud computing that an organization needs to truly get the best of both worlds.

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

When the Edge Mesh is designed into a print solution, organizations can have a true enterprise printing experience with the budget of a small to medium-sized business. High security, near-zero latency, and high scalability are all built into the Edge Mesh's architecture. With all of these features working together to replace the server as we know it, administrators can rest easy knowing they don't have to worry about 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 infrastructure 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, and 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 edge node selection (part of the printing process) toward themselves to further improve high-availability printing.



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 scalable. As a new host is added to the printing environment, it can commit some of its resources to the Edge Mesh.

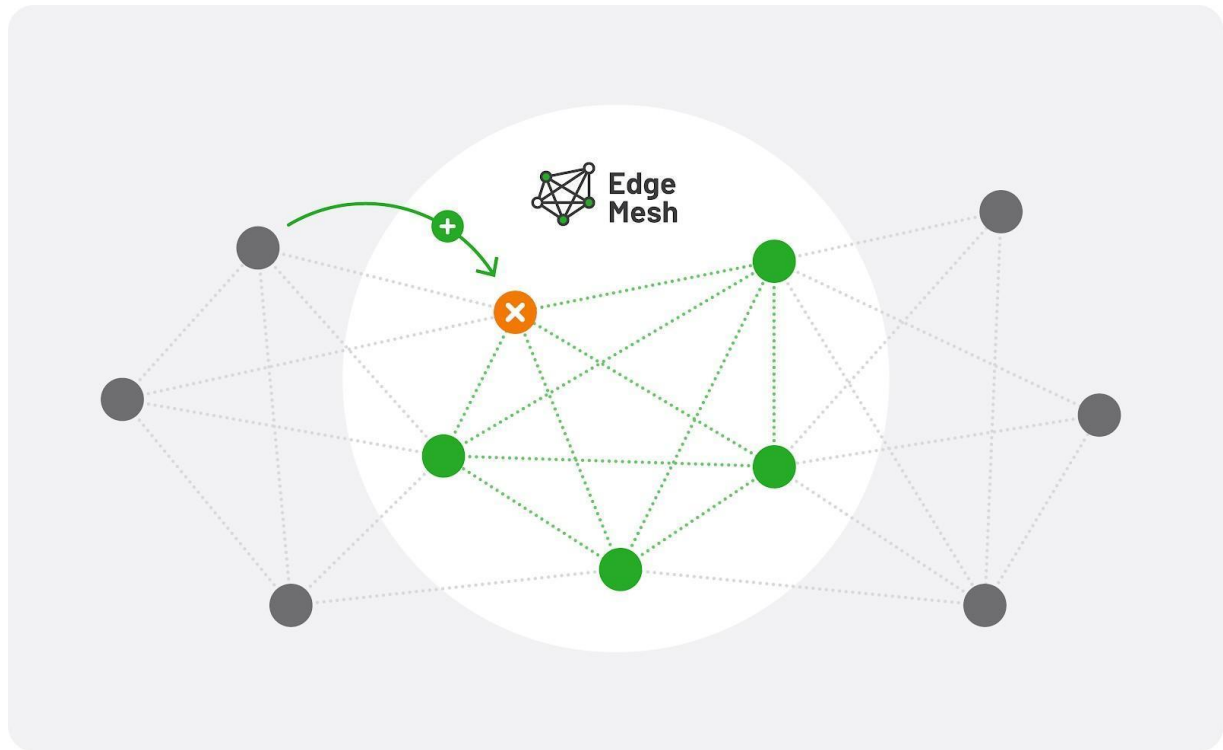
As your printing environment grows, so does the number of hosts capable of handling the processed print jobs. If a single edge node fails, another node in the Edge Mesh will take its place. As long as one node is online, job processing continues.

The Edge Mesh's high availability doesn't end with hardware failover. The mesh also applies high availability to the jobs it processes. Print jobs are triple-protected from the failure of any single node since, on job 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 printing operations. For example, nodes within the Edge Mesh communicate with one another to determine where the job is and where it needs to go to get printed.

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 to the required printer. With these mechanisms in place, it takes a very significant network disruption to stop printing (and, by then, you've probably got other things to worry about...).



As more computers are added to your network, so does the number of nodes capable of handling the processed jobs.

Additionally, network administrators don't have to worry about scheduling maintenance windows to update server software when issues or inefficiencies are discovered.

In a traditional server environment, server downtime needs to be scheduled for upgrades to take place. The risk of losing print functions during the upgrade just had to be accepted. With the Edge Mesh, the edge nodes individually communicate with the cloud for up-to-date network routing information. This means that the nodes can patch themselves in a way that maintains enterprise availability for the entire print infrastructure. No downtime. No data loss.

Maximum security

The 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 that in an Edge Mesh, no organization outside of your own ever handles your documents.



Secure local printing

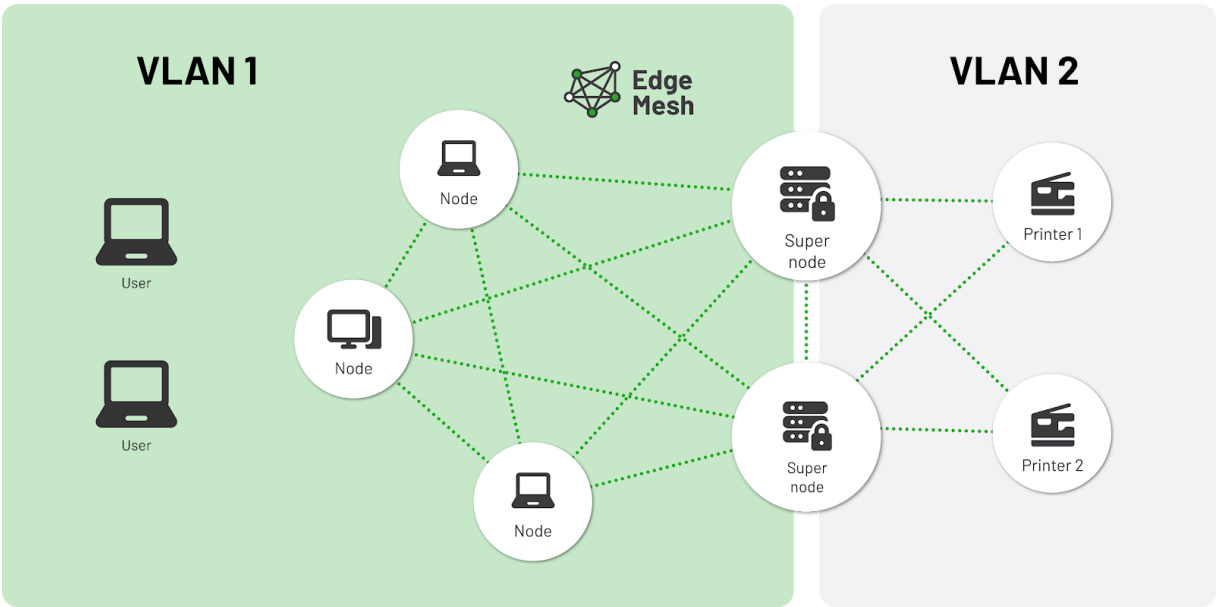
Of course, just “being local” does not necessarily mean “being secure”.

<p>“From a security perspective, Zero-Trust can better track and block external attackers, while limiting security breaches resulting from internal human error.”</p> <p>Defense Innovation Board, 2019</p>	<p>“Average Size of a Data Breach: 25,000+ Records”</p> <p>“Average Cost per-record: \$242.00 USD”</p> <p>IBM, 2020</p>
---	---

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, the cloud governs access to the mesh network, which challenges new hosts to verify access to the Edge Mesh.

Since the cloud is the source of all the Edge Mesh’s initial configurations, trust is built into the mesh without the potential of compromising security.

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 (for example, you might have a complex password on your WIFI, but connecting to the network using a printer’s ethernet cable might make this security moot!).





PaperCut’s Edge Mesh helps with network security, too. The mesh contains an intelligent routing algorithm so jobs can “hop” via nodes in the Edge Mesh to navigate to other network segments.

Print jobs can flow by traversing a boundary node connecting the VLANs. One common method to enable this is to configure a Virtual Machine with edge node software, granting it access to both VLANs.

In a similar methodology used with Kerberos, the cloud becomes the manager of the relationships between computers and printers 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.

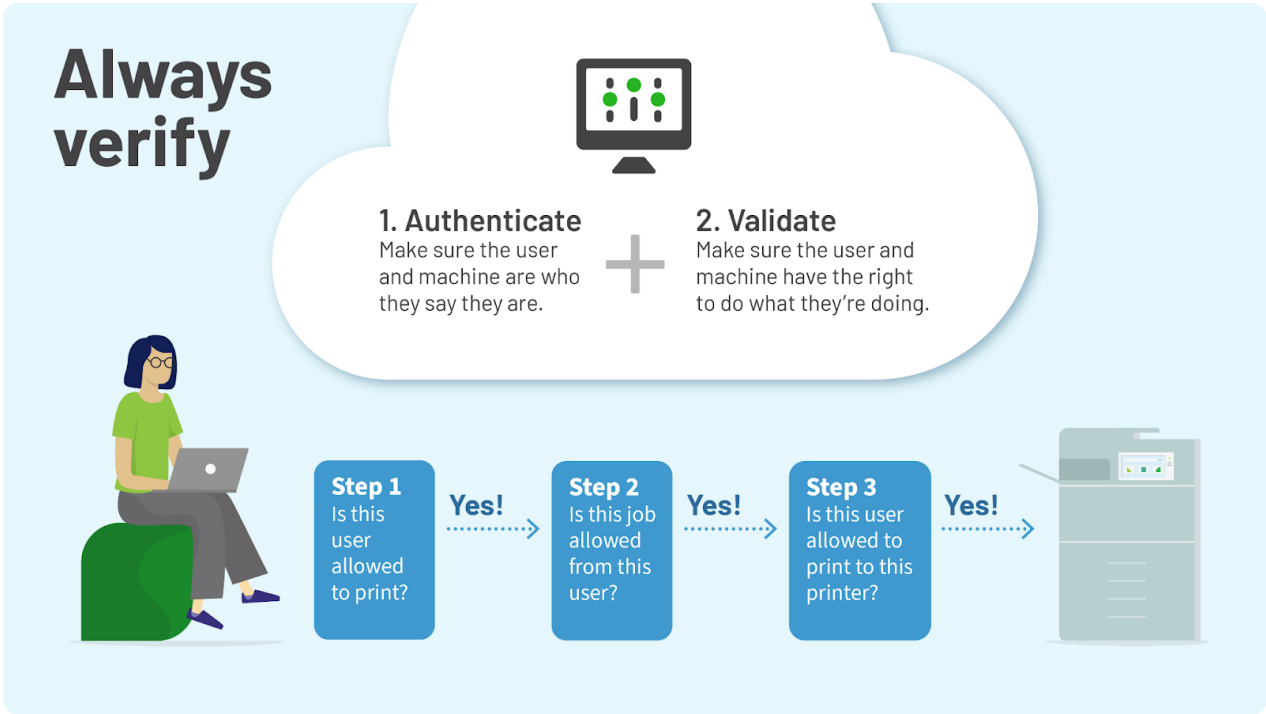
Zero Trust networking

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, 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 from the node that has left the mesh.

Principles of Edge Mesh zero trust





Every step in submitting and printing 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 encrypts the data using a three-part key provided by the cloud and stores 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 it 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 key 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 environments. Zero Trust guards against all access-based compromises and protects data at every point within the Edge Mesh. The Edge Mesh makes zero-trust accessible to any administrator because it's built into the way the mesh works.

Looking to the future

The Edge Mesh has the potential to shift how printers work in networks. Mesh technology brings to networks the practically-zero latency and high security of local network resources, along with the massive scalability and high availability of using 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, the functionality doesn't come at the cost that's traditionally associated with high performance. Edge Mesh stands as a democratizing shift in the structure of our networks. No longer is the barrier to entry of high availability insurmountable for the small-to-medium-sized business.

For organizations of every size, the Edge Mesh drives performance and security and is a significant first step in spreading the use of high-security, zero-trust networking to all organizations.

PaperCut believes that the 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 PaperCut Hive.



Works cited

McCue, Jim “[ILTA 2019 Technology 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