



Equinix Digital Support & Success Guide

How to get the most out of Equinix Fabric,
Equinix Metal and Network Edge.




Customer success is at the heart of what we do.

At Equinix, customers trust us to operate their critical digital infrastructure. In addition to a stable, world-class network and a bare metal automation platform that is second to none, true success often comes down to people, process, and procedure. This document is intended to provide details about how our team is structured, how we operate, and how we can help you succeed with our digital products from first touch through emergency incident.

If you have questions or feedback, I'd love to hear from you.



Jeremy Brooks
Sr. Director, Digital Support and Success
Equinix



Contents

This Document Covers Three Products	4
Documentation & Resources	5
SLA's & SLO's	6
Shared Responsibility Model	7
Status Page & Alerts	8
Support Channels & Contacts	9
Special Consideration: Equinix Metal Hardware	10
Special Consideration: Network Edge Resiliency & Redundancy	11

Maintenance & Incident Management	12
Leadership Contacts	13
Equinix Fabric FAQs	14
Equinix Metal FAQs	15
Network Edge FAQs	16



This Document Covers Three Products



Equinix Fabric®

Directly connect your infrastructure to thousands of destinations, from major public clouds to network service providers and enterprises.



Equinix Metal®

Equinix Metal is a bare metal as a service that offers high-performance, on-demand, and automated infrastructure. Deploy globally in minutes.






Network Edge

Deploy virtual network services close to users, clouds and networks—the digital edge—in minutes without additional hardware.

Documentation & Resources

Documentation is central to how we deliver an excellent experience to customers.

Product	Product Documentation	API Docs	Guides & Tools	Support Email
 Equinix Fabric	docs.equinix.com	developer.equinix.com	docs.equinix.com	support@equinix.com
 Equinix Metal	deploy.equinix.com	deploy.equinix.com	deploy.equinix.com	support@equinixmetal.com
 Network Edge	docs.equinix.com	developer.equinix.com	docs.equinix.com	support@equinix.com

To view our roadmap or request features for any of these products, please visit feedback.equinixmetal.com.



SLA's & SLO's

SLA's and SLO's help align us with our customers. While outages do happen, our shared goal is to help customers meet their high availability and redundancy goals through careful planning and best practices.



SERVICE LEVEL AGREEMENTS

Our SLA's relating to hardware, network and platform services can be found in our Product Policies, linked below. We've also provided answers to some frequently asked operational questions (for instance about hardware failure in Metal, high availability in Network Edge, and redundancy in Equinix Fabric) later in this document.

SERVICE LEVEL OBJECTIVES

Equinix does not currently offer a guaranteed response time on tickets or chats. However, our historical average response time across channels is under 30 minutes.

PRODUCT POLICIES

To view our SLA's, please review each of our product policy documents:














- [Equinix Fabric](#)
- [Equinix Metal](#)
- [Network Edge](#)



Shared Responsibility Model

 **Responsibility retained by customer**

 **Responsibility transfers to Equinix**

	DCaaS Data Center as a Service	NaaS Network as a Service	BMaas Bare Metal as a Service	IaaS Infrastructure as a Service
Application, Runtime, APIs and Middleware Availability and Security				
Identities and role management in Equinix Portals				
Internal host networking				
Guest OS				
Virtualization				
Physical host				
Physical network and configuration				
Data Center				
Business Contact Information, metrics and metering data, service logs				
Equinix Business Applications and Services, Portals and APIs				

View the full Shared Responsibility Model documentation on our website.



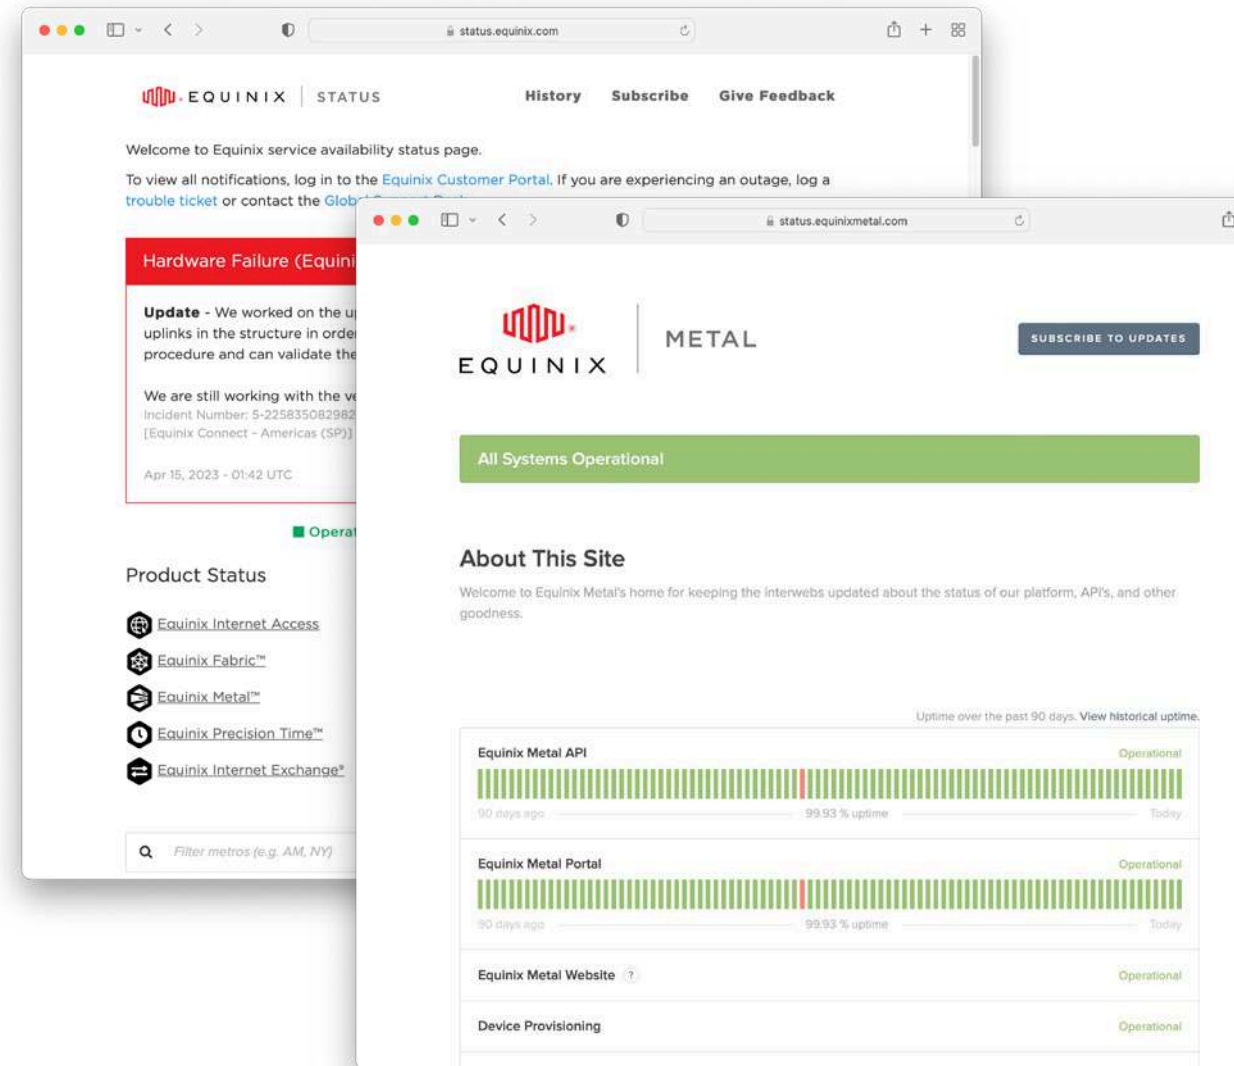
Status Pages & Alerts

Equinix proactively communicates all platform impacting events as soon as they are known. This includes reactive events (outages, incidents) as well as proactive maintenance events that will be service impacting.

Customers are encouraged to subscribe to our [Equinix status page](#) for timely alerts of high severity incidents, as well as followup communications (Reasons for Outage, or RFO's). If you're an Equinix Metal customer, we also encourage subscribing to the [Equinix Metal status page](#) for maintenance notifications and lower-level incident notifications.

TYPES OF UPDATES POSTED

- During every impactful incident, the Equinix team provides all live updates to this status page in order to keep all customers aware of any occurrences in the environment
- Planned Monthly or Emergency Maintenance notifications



Support Channels & Contacts

Equinix Metal:

- Open a ticket through the [Metal Console](#)
- A live chat on [deploy.equinix.com](#) or [console.equinix.com](#)
- Email support@equinixmetal.com

Equinix Fabric & Network Edge:

- Tickets can be opened in the [Equinix Fabric portal](#)
- Email support@equinix.com.

For high severity issues, please open a ticket and include a description of the impact. In all cases, please provide as many relevant details as possible, including error messages, traceroutes, MTR's, etc. Thank you!

Special Consideration: Equinix Metal Hardware

As a provider of dedicated, physical infrastructure, the stability of our hardware directly impacts our customer experience.

While we cannot promise that hardware never fails, here is what we can deliver:

- **Quality**—We invest in premium hardware and related datacenter components.
- **Best Practices**—We help you architect for the right level of resiliency on our platform.
- **Replacements**—We keep appropriate replacement capacity on hand for top configs.
- **Dedicated Spares**—We sell dedicated sparing for critical or customized infrastructure.

When it comes to hardware failure, **our SLA is simple**—we will make available new hardware to replace your instance. If your monitoring detects a problem with hardware, please validate the issues you are seeing and then engage with our support team. If a replacement is required, here is our SLO (e.g. how it works):

- For on-demand Standard Gen3 hardware, we stock capacity available via our API or portal. Simply find the closest match to your existing hardware and deploy. Let our support team know about the failed hardware and they'll take care of fixing it up!
- For reserved hardware (Standard Gen3 or Workload Optimized) we keep facility spares on hand to swap out for any failed hardware. Please coordinate with our support team for replacement, which can take up to 1 business day.
- If you have contracted spares, our support team can reserve them to your project 24/7.

PLEASE NOTE:

- Clients focused on resiliency often deploy an “N+1” architecture for maximum redundancy and control.
- We do not schedule drive repairs etc as those take far longer than the few minutes to get a new machine. For custom machines we require spares for this purpose.

Special Consideration:

Network Edge Resiliency & Redundancy

Designing solutions with resiliency is one of the most critical aspects of network and edge architecture. While there is no correct answer to how much resiliency is needed, there are best practices, suggestions for different use cases, and some specific services and features that Network Edge offers.

The underlying Network Functions Virtualization (NFV) platform that provides the infrastructure for Network Edge is inherently fault-tolerant from a single virtual instance standpoint. Still, you must design high availability into the overall solution to achieve the maximum redundancy possible, including multi-metro architecture for many use cases.

For full details and best practices, please review the [Network Edge documentation](#).

Maintenance & Incident Management

Postmortem process for high severity incidents reported on status.equinix.com

WHEN DOES EQUINIX ISSUE AN RFO?

Reason for Outage (RFO) reports are created at the request of any customer-facing team (Customer Success, Technical Account Manager, or Sales). This ensures the work of creating a report adds value.

WHAT IS THE TIMELINE FOR REQUESTING A RFO?

A request for RFO may be received anywhere from during an incident up to approximately 72 hours after an incident has occurred. Any RFO request received later than 72 hours post-incident will receive best effort, but it may be more difficult for teams to track down relevant information.


WHAT IS THE TIMELINE FOR RECEIVING A RFO?

Equinix's RFO timeline commitment is 72 hours from the time of RFO request. The Equinix Operations team will expedite as possible for urgent requests, and preliminary RFOs may be created as needed for incidents which are not yet fully understood.



Leadership Contacts

The entire Equinix team is committed to creating and operating services that you can trust for your most critical infrastructure. Our Digital Success team welcomes your feedback and about what we're doing well, and how we can improve.



Jeremy Brooks

Sr. Director, Digital Support & Success

jebrooks@equinix.com



Josh Kelly

Sr. Manager, Technical Account Manager

jokelly@equinix.com



Jamie Law

Sr. Manager, Sr. Customer Engineering

jamlaw@equinix.com



Pim Jeurink

Customer Engineering, Europe / Middle East / Africa

pim.jeurink@eu.equinix.com



Jemarie Camba

Customer Engineering, Asia-Pacific

jcamba@equinix.com



Bryce Adams

Customer Engineering, Americas East

badams@equinix.com



Gajanan Sane

Customer Engineering, Americas West

gsane@equinix.com



Equinix Fabric FAQs

Q: How is redundancy managed on Equinix?

A: Equinix Fabric is operated as a fully redundant platform with [Primary and Secondary chassis](#) and customers are able to choose the right level of redundancy in their [architecture](#). Equinix builds redundancy into our network by deploying multiple high-capacity circuits between metros, resulting in a highly resilient network which can accommodate multiple failures.

Q: What is the difference between a remote port and a remote connection?

A: A remote port is used when you are not co-located in an Equinix IBX and you use an NSP to connect from your location to Equinix Fabric. See the Port Types page for more information. A virtual connection is used to connect to a provider, partner, or even yourself in an Equinix Fabric market (such as Chicago) when you are in a different Equinix Fabric market (such as New York).

Q: How soon can I use my newly installed Equinix Fabric Port?

A: As soon as it is provisioned, your port should be usable. If you have followed the troubleshooting section in our Getting Started page and still have issues with your new port, please open a ticket via the Fabric portal or email support@equinix.com.

Q: What are the Equinix Fabric SLAs for latency, up-time, and provisioning?

A: The [Equinix Fabric Product and Policy Document](#) is a guiding contractual template explaining general Equinix definitions, policies, SLAs, and the Terms and Conditions of doing business with Equinix. The Equinix Fabric Product and Policy document is a guiding document, and not a contract.

Q: How do I connect my router to Equinix Fabric?

A: The Equinix Fabric port and cross connect to your cage/cabinet/NSP is included as part of the Fabric port. All that is left to arrange is the patching within the cabinet to your equipment or NSP, which your engineers can complete. For cage or cabinet connections, you can also [Order Smart Hands](#) to have Equinix personnel complete the connection on your behalf.

Equinix Metal FAQs

Q: What happens to hard drives that get replaced on a server when a disk fails?

A: Hardware is recycled following a standard policy that includes a wipe of the disk and/or physical destruction.

Q: Are contracted spares racked, enrolled and powered on?

A: Your committed spares are racked, enrolled, and reserved to your project, in a powered on state. When you have a hardware failure and contact support, they will help move that spare to a provisionable state.

Q: What happens if a component fails and I don't want to migrate to a new server due to rebuilding and data migration challenges?

A: Equinix Metal will attempt to repair the hardware on a best effort basis, however it is best to architect your deployment for redundancy.

Q: What happens to my data on my current server when I move to a new one due to hardware issues?

A: The data will be wiped when the hardware goes through a deprovisioning process.

Q: What is the best way to maintain high availability in the case of hardware failure?

A: Having the right architecture is a critical part of meeting high availability and disaster recovery scenarios. In addition, dedicated sparing for mission-critical use cases or fragile application environments is recommended.

Network Edge FAQs

Can I customize the Network Edge software or hardware in the stack?

A: The Network Edge stack includes the CPU, memory, storage and other elements of the service. You can't customize or make changes to the stack that is provided as part of the Network Edge package. Additionally, the raw compute specs (CPU, memory, storage) are pre-determined per device selected.

Does Equinix manage the devices I deploy into Network Edge?

A: No, Equinix manages the Network Functions Virtualization Infrastructure, NFVi, but not the VNFs themselves. We do have many partners we can refer to for management services.

How are redundant devices deployed?

A: Each Network Edge metro includes two planes with separate compute and top-of-rack infrastructure. These two planes are used for primary and secondary deployments. Learn more in our [Architecting For Resiliency docs](#).

Is my virtual port automatically redundant?

A: Yes, connectivity between the Network Edge infrastructure, Equinix Fabric and Equinix Internet Access (EIA) is fully redundant providing uninterrupted connectivity in the event of an interconnection failure.

Can my virtual devices connect each other across metros?

A: Yes! We have “any to any” connectivity. Devices in multiple Metros or within a Metro can be connected using our [Device Linking feature](#) or with [EVP-LAN](#).

