eBook

trustgrid

Connecting Distributed Applications

www.trustgrid.io

Contents

01 **Overview**

trustgrid

Delivering applications from the cloud or data center has become more complex than ever.

Applications have shifted from running in environments where data and compute reside behind the same firewall to architectures dependent upon customer-controlled data, microservice APIs, and on-premise deployments integrated to cloud services.

These evolutions have caused SaaS application providers to re-think their infrastructure as they seek to minimize operational friction, optimize customer experience, and remove barriers to product innovation.

In this eBook we will show you how Trustgrid is enabling SaaS application providers to build distributed applications when they are delivered from the cloud, but require connectivity to, or deployments in, on-premise environments.

02 What is a distributed application?

For many application providers the cloud is simply not enough. These providers can centralize some, but not all, of their systems in a single cloud environment. They find it necessary to deploy in disparate geographies, across hundreds or thousands of sites, in customer environments, and in various public and private clouds.

Why distributed applications are needed	
Hardware Adjacency	Control systems that interface on non-routable protocols like an MRI machine, fuel pump, or door security latch.
Legacy Integration	Centralized applications require integration to legacy systems in the cloud or in private data centers. This could be cloud apps connecting to customer-controlled data sources or mainframes being leveraged by cloud services.
Remote Survivability	Systems must continue to function even when the internet is down such as retail point-of-sale terminals and bank ATMs.
Cost Containment	The costs of public cloud can require optimization through distribution of key components in lower cost environments.
Third Party Integration	More and more of today's applications leverage third party services through API integration.

These applications frequently centralize most of their services, but still rely on a number of components deployed outside of those environments. Often, they are dependent on both third party vendors and customers. This lack of control can introduce challenges in deployment, maintaining connectivity due to internet issues, and even support. In addition to a cloud DevOps team, they almost always have a NOC responsible for on-prem connections and deployments.

In short, distributed apps have many more dependencies than traditional centralized applications. These architectures combine all of the challenges of cloud with the difficulties of on-prem deployment.

03 How do distributed applications function?

Software vendors hosting SaaS applications that require connectivity into other systems or customer environments will require a networking component. This network may be connecting cloud-hosted apps to on-premise data sources, or inversely allow applications running on-premise to have access to services in the cloud.

Historically, IPSec VPNs have been used to create these tunnels. However, legacy VPN solutions come with limitations that can inhibit the ability for software providers to build additional features or increase the efficiency of its operations.

The Challenges of Connecting Distributed Applications

Initially, a software vendor's understanding of the challenges surrounding distributed applications is focused on the most important one or two issues preventing success. As development progresses and deployments scale, the extent of the challenges come into greater focus.

Commonly, deployment challenges surface when applications that are dependent on appliances deployed by an end user in their data center want to evolve into cloud-dependent services with APIs.

Additionally, SaaS product and engineering leaders may look at tools like VPC peering and VPN gateways, but neither of these solve the scalability challenges of managing hundreds or thousands of connections and VPNs to customer environments.

As a SaaS solution evolves, so do its architectural requirements.

Often, connectivity solutions that were viewed as acceptable yesterday no longer work for the roadmap and profitability requirements of tomorrow. Containerized services or automation of support are just a couple of features that may be needed in the future, but are prevented from implementation due to ill-informed architectural choices early in the project.

SaaS vendors need a more automated connectivity solution that does not require the management of hundreds (or thousands) of unique VPN configurations. The Trustgrid platform solves many of the issues encountered by application providers as they attempt to connect, deliver and scale applications from the cloud to the edge.

04 The Platform

The integration of distributed environments requires deployment, management and support features that typical hardware-centric networking solutions, and even most modern software-defined networking technologies, are not designed to handle.

The Trustgrid platform integrates a unique combination of SD-WAN, remote access and edge computing into a single application platform. The platform supports the challenges of securely connecting and delivering applications to any system or device, regardless of where it is located.

The Trustgrid platform's multi-disciplinary approach to connectivity and application delivery combines managed services and automation tools with a distributed application platform.

When combined, these tools allow DevOps and IT administrators to build an extensible fabric of network security and enables application deployment across any environment.

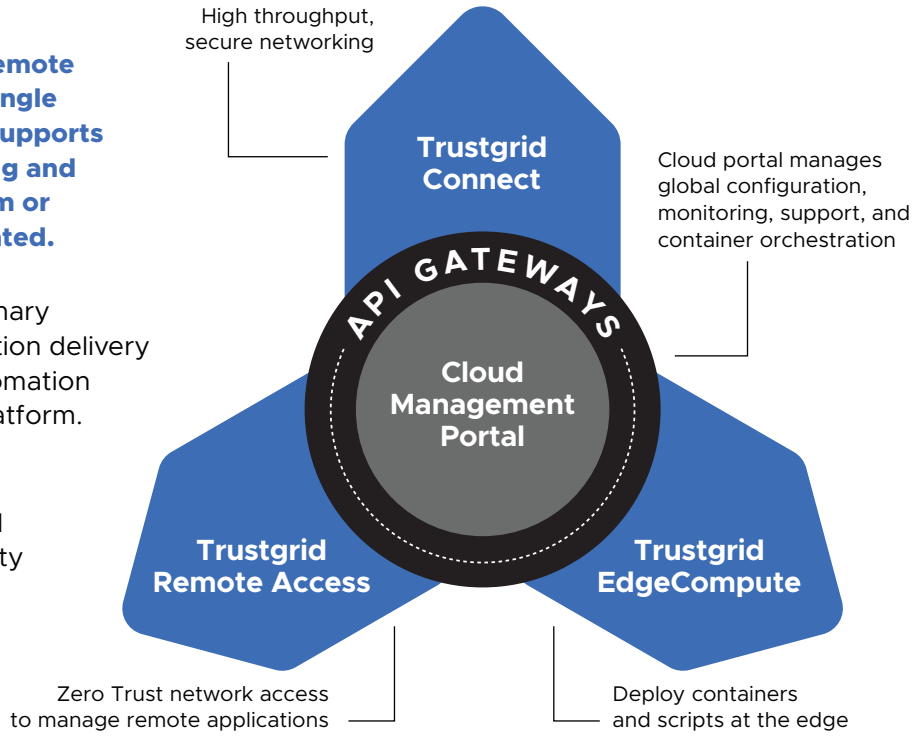

The Platform consists of 3 products; Connect, EdgeCompute and Remote Access.

Features of the platform include:

Interfaces

- Cloud Portal
- Management APIs

Control Plane Services

- Identity Based Access Control
- Virtual Networking Overlay
- Cloud PKI
- Cloud Monitoring
- Software Updates
- Centralized Configuration
- Remote Management Tools
- Container Registry
- Centralized Reporting

Data Plane End Points

- Cloud Nodes
- Hardware Nodes
- Virtual Appliance Nodes
- Open Source Agents
- Agentless Portal

4A **Interfaces**

The platform's control is centered around a cloud-delivered management portal hosted in AWS. As an API-first platform serving software developers and DevOps teams, all functionality is built first as an API before a UI is applied. This allows for the full suite of platform capabilities to be leveraged via code, inserted into other applications or even built upon according to our customer's own needs.

The cloud management portal houses all configuration, monitoring and support tools, and 3rd party integration services of the platform. It serves as the control plane for all network services.

4B **Control Plane Services**

Managed from the cloud management portal, the control plane services are the platform's core components and features.

Identity-Based Access Control

Build fine-grained, zero trust access policies based on user or device attributes such as roles, groups, geography, IP ranges, and more. The access control features can be built using existing IdPs (Azure AD, Okta, Google) or built independently within the Trustgrid system.

Virtual Networking Overlay

Overcome the network address translation (NAT) issues that arise when connecting to many different organizations or networks by creating a virtual network overlay that maps similar internal IP addresses with custom labels that you set. Our private data plane architecture ensures that no customer traffic is ever visible to Trustgrid.

Cloud PKI

Ensure data security with our proprietary cloud PKI to authenticate and encrypt user and device traffic. Trustgrid offers options for customers to hold their own keys or let Trustgrid manage them.

Cloud Monitoring

Create a virtual network operations center with single pane of glass visibility, performance monitoring, and automated status alerts can send notifications via Slack, PagerDuty, OpsGenie, SMS or other communications tools.

Software and OS Repo

Ensure security, compliance, and always have the latest features with automated patching and updating of all connected systems.

Centralized Configuration Management

Automate version control and configuration backups. Deploy network nodes and push configuration updates from the cloud management portal.

Remote Management

Remotely control the entire Trustgrid network, including the ability to manage containers and adjacent customer devices at the edge.

Container Registry

Centrally store Docker containers for automated deployment across any Trustgrid node.

Centralized Reporting

Simplify compliance reporting with single source of truth logging of access, performance events and other telemetry. APIs allow 3rd party SIEM and infrastructure monitoring tools to pull in all traffic and change logs.

4C **Data Plane End Points**

The Trustgrid platform is natively "multi-premise" with the ability to connect and deliver applications across all environments using our suite of endpoints.

Cloud Nodes

Public or private clouds can leverage cloud-optimized virtual appliances to build connections to cloud applications or data sources. These nodes can serve as gateways connecting a centralized application to hundreds (or thousands) of other Trustgrid enabled endpoints. Trustgrid supports cloud nodes in AWS, Azure, Google Cloud, and Oracle Cloud.

Hardware Nodes

The Trustgrid endpoint software can be deployed on any x86 hardware. Customers may source their own hardware or work with a Trustgrid partner.

Virtual Appliance Nodes

Similar to the hardware nodes, Trustgrid endpoints can be deployed as virtual appliances on a customer's existing hypervisor infrastructure. Virtual Appliance nodes have the ability to run any of the products in the Trustgrid platform.

Open Source Agents

The Trustgrid platforms brings IDP integration, granular security controls, and centralized management to open source agents from OpenVPN and Wireguard™.

Agentless Portal

As the preferred method of delivering Trustgrid Remote Access, our Agentless Portal provides zero trust network access to remote users from a web portal. Leveraging an organization's identity provider, authorized users have private access to applications without requiring a device agent.

05 The Products

Each of the products in the Trustgrid Platform work together seamlessly and are managed from a common management portal.

5A Trustgrid Connect

Trustgrid Connect is a network-as-a-service delivering next-gen SD-WAN capabilities. It is designed to meet the challenges of application providers who require cloud-to-on-premise and multi-cloud networking.

Specifically designed for SaaS applications that must connect to hundreds or thousands of customer, partner, or other diverse IT environments, Trustgrid Connect is an alternative to site-to-site VPNs and MPLS to provide a cloud-delivered WAN, optimized for ease of management.

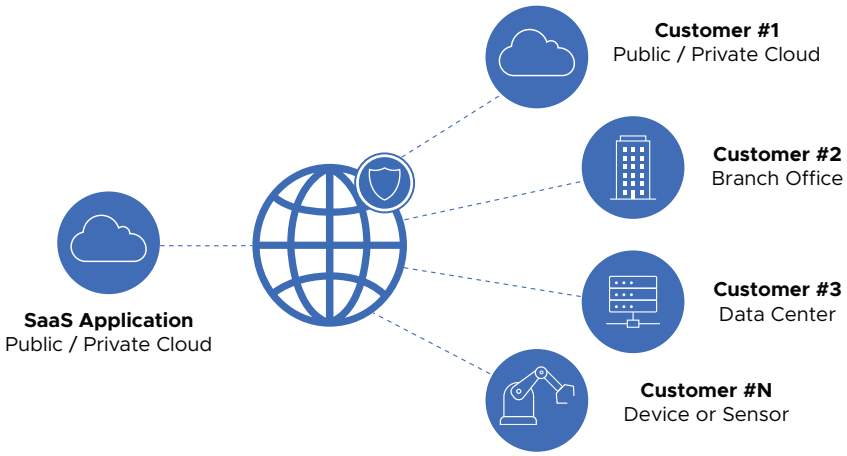

Trustgrid Connect builds a multi-tenant network fabric between a cloud application and an any number of edge environments.

This WAN supports both mesh and hub-and-spoke architectures to support multi-cloud, hybrid cloud (cloud to on-prem), and on-prem to on-prem use cases. Trustgrid can offer IP SLA to select the best path for traffic, as well as QoS at the end points to prioritize latency sensitive data. Segmentation by customer is default behavior and can integrate to VRFs and VLANs in data centers and cloud environments.

Application providers gain global visibility, control and support capabilities from a single-pane of glass that logically separates each network from each other. This multi-tenant connectivity allows for an application provider to support the entire network as if it is another cloud delivered service.

Features include:

- Layer 3 / 4 networking
- Zero trust network architecture
- Separate control plane and data plane
- Certificate-based authentication
- Continuous patching and updating
- Automated failover and disaster recovery
- Supports all cloud and on-premise environments
- Simplified network address translation management
- 1-touch deployments with little to no firewall reconfiguration

Trustgrid Connect works seamlessly with Trustgrid EdgeCompute or Trustgrid Remote Access.

5B Trustgrid EdgeCompute

Trustgrid EdgeCompute adds to the capabilities of Trustgrid Connect to provide a distributed computing platform for deploying and supporting applications at the edge.

Integrating networking features with a containerized application platform, Trustgrid EdgeCompute overcomes the challenges of edge computing by delivering infrastructure, services, APIs and software lifecycle management tools to applications running in on-premise locations.

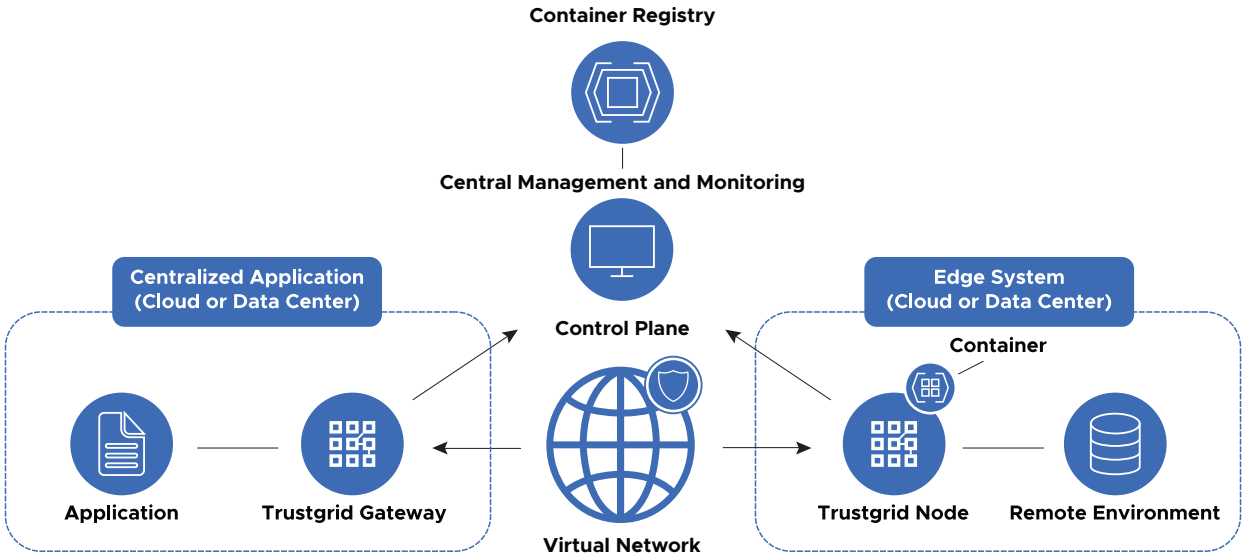

The platform runs in all public and private cloud environments with plug-and-play deployments that eliminate the need for on-site networking or container expertise.

For application providers building applications that require low latency, local data processing, or need to meet data residency requirements, EdgeCompute creates a seamless distributed application delivery environment that allows edge appliances and services to be managed in the same way as a cloud service.

Unlike AWS Outposts which eliminates the ability to expand into multi-cloud architectures or Google Anthos which typically requires teams of advanced engineers to support its complex Kubernetes-based architecture, Trustgrid EdgeCompute is designed to run cross platform and be maintained by as little as a single staff member.

Features include:

- Edge computing with seamless networking
- Cloud managed container repository
- Support for Docker containers and KVM virtual appliances at the edge
- CI/CD integration for automated patching and updating of remote systems
- Ability to build and maintain APIs for any data source
- Supports 1000s of remote services from single pane of glass
- Run 3rd party security solutions on at any edge location

Trustgrid EdgeCompute works seamlessly with Trustgrid Connect or Trustgrid Remote Access.

5C Trustgrid Remote Access

Trustgrid Remote Access is designed to provide secure, granular access to remote application components for patching, troubleshooting and support.

Legacy VPNs or remote desktop support tools lack the granular access management controls needed for sensitive application components and can be easily exploited via stolen credentials and session hijacking. Extending remote access to 3rd parties or vendors for support can introduce even more risk.

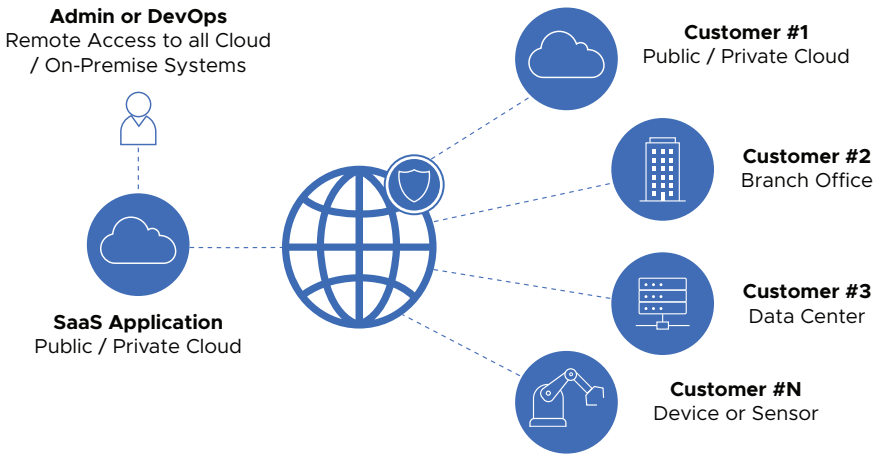

Trustgrid Remote Access provides zero trust network access (ZTNA) for software administrators and DevOps teams supporting distributed application deployments.

Trustgrid Remote Access supports applications running in any cloud or on-premise environment by using an agentless web interface to grant access across a range of applications and devices. Access policies can be custom configured within the tool utilizing existing roles, groups and permissions from 3rd party identity providers. If an agentless approach is not desired, open source agents such as OpenVPN or Wireguard are also supported.

Features include:

- Agentless zero trust network access
- Supports access to apps on all popular operating systems
- Integrates with your existing IdP (Azure AD, Okta, Google, and more)
- Maintains auditable logs of all application access
- Provides secure access to remote container deployments
- Remotely supports systems behind firewalls you don't control

Trustgrid Remote Access works seamlessly with Trustgrid Connect or Trustgrid EdgeCompute.

06 Deployment

The Trustgrid platform deploys in two phases. First, in our customers' environments and then into the end user environments. In the first phase, Trustgrid's professional services teams assist with deploying Gateway Nodes in your public and private cloud infrastructure as well as:

- **Configuring portal and API access**
- **Network design**
- **Integration to monitoring/IDP/SIEM and other tools**
- **Customization of end user documentation and on-boarding materials**

In the second phase, Trustgrid assists our customers with on-boarding new end users or migrating existing end users from current connectivity or deployment methods. For new end users Trustgrid provides customized documentation and due diligence materials to quickly overcome any objections about introducing a third party to the system.

The end user is asked to complete an encrypted web form with their desired appliance configuration (IP, DNS, failover method). Appliances are shipped (or emailed a link for virtual appliance download) directly in a ready-to-install package. All order status updates are contained in the Trustgrid portal for easy access by customer success teams.

When a customer seeks to migrate away from an existing solution Trustgrid can provide a professional services engagement to do so. This is frequently the case during end-of-life/support instances where many devices must be rapidly replaced. Trustgrid takes the lead in these engagements and works directly with end users to minimize the burden on our customers. The process for each end user is nearly identical to the 'new end user' process above but is centrally managed by the Trustgrid team.

6A Appliance Selection

Trustgrid network nodes are software-defined and allow customers to choose what works best for their situation.

Hardware

Trustgrid software is imaged to a plug-and-play, secure network appliance.

Virtual Appliance

Create a node without dedicated hardware in any environment supporting VMWare or Hyper-V.

Cloud

Build a cloud network endpoint directly from AWS, Azure, Google Cloud or Oracle Cloud.

6B Availability and Disaster Recovery

Depending on availability requirements, network nodes can be paired at any location and across multiple data centers to ensure an always-on connection. When configured for maximum redundancy, Trustgrid offers a 99.99% uptime SLA guarantee.

Single Node

The most basic configuration used to establish a network endpoint.

HA + Single Node Disaster Recovery

A redundant pair of nodes is used for a network endpoint with another node placed at a secondary data center to provide always-on connectivity should an entire location fail.

High Availability (HA)

A redundant pair of nodes configured to enable automatic failover should a problem arise at a single location.

HA + High Availability Disaster Recovery

A redundant pair of nodes is placed in two different data centers to provide the highest levels of failover and disaster recovery for mission critical application connectivity.

6C Network Configuration

Selecting the number of interfaces a node will have depends on your preference for public/private IP addresses and the locations of firewalls surrounding the Trustgrid endpoint.

One Interface

A single network interface (one arm, on a stick) is configured on the local network, behind the firewall.

Two (or more) Interfaces

Two interfaces are configured to accommodate placing the node in a WAN / LAN or DMZ (behind firewall) / LAN configuration. These nodes may be placed with the WAN interface on the public internet as well.

Once the configured appliance arrives, a user simply plugs in the device (hardware only) or follows the instructions for virtual appliance installation.

For endpoints needing to be installed behind a firewall, a handful of clearly defined IP ranges and ports will need to be allowed outbound for connectivity to our customer's Gateway Nodes and the Trustgrid control plane. The Trustgrid appliance will then authenticate, download any necessary updates, and be ready for use.

07 Management

The Trustgrid platform is continuously improving with the ability to easily push updates to all connected nodes

The Trustgrid platform consists of two primary elements. The cloud management portal and endpoints referred to as nodes. The cloud management portal serves as the central interface for the platform and contains all of the necessary management controls, integration tools, monitoring, and support functions.

Integration to an identity provider (IdP) governs access to the portal and access can be configured to share with any internal or external stakeholders. Fine-grained and customizable permissions can grant or restrict access within the portal so that administrator rights are differentiated from those who need simple visibility into performance metrics.

All management of nodes occurs from the management portal or the API. Configuration changes, patches, updates and troubleshooting can be done individually per node or grouped for batch changes to streamline network operations.

The platform operates under a CI/CD (continuous integration and delivery) methodology that pushes patches and security updates frequently (about once a month). This ensures that the system is always in compliance, always running the latest security patches, and providing users with the latest feature enhancements.

Updates are pushed to nodes during maintenance windows set by the user, pushed to all connected nodes using a blue/green methodology, and can easily be reverted to a previous state should a problem occur.

The management portal functions as the control plane of the network while the nodes (or gateways) serve as the data plane. The data plane passes all traffic between the nodes and operates separately from the control plane to ensure that should the cloud management portal experience problems, the data will continue to pass between connected nodes. Additionally, should a data plane endpoint experience a problem, the control plane will have the ability to remotely troubleshoot the issue.

08 Support

The management portal allows for all connected nodes to be remotely monitored, diagnosed and supported. The portal has visibility to system meta-data (IP addresses, events, packet/byte counters, etc) created by the nodes, but no ability to see the actual data being transmitted. All meta data is logged and can be referenced for retroactive investigation should an event occur.

Notifications can be set to alert administrators (via Slack, PagerDuty, OpsGenie, SMS and more) of activity that exceeds predefined thresholds.

When a node experiences problems that require investigation, nodes can be accessed via cloud terminal, event logs can be pulled, and a node can be remotely restarted or updated if needed.

When a performance issue occurs, Tier 1 support inquiries are typically handled by the Trustgrid customer's support team and escalated to a Trustgrid engineer if a solution is not immediately discovered.

This architecture also means that traffic between nodes is private and not visible to any entity that has not been specifically authenticated and authorized. Trustgrid's support team offers 24/7/365 support for incident response.

09 Use Cases

The Trustgrid platform is designed to help applications connect and deliver services at the edge.

The most common Trustgrid use cases.

9A Cloud-hosted application connecting to customer data center

Some applications that run in public cloud environments must leverage data that resides in a customer's cloud or data center environment.

Whether due to difficulties migrating the data to the cloud, customer security concerns, or legally-required data residency issues there are numerous instances where a SaaS application must have real-time, bi-directional access to data behind a firewall they don't control.

In these situations, Trustgrid can be used to establish cloud gateways in front of a customer's cloud application.

This gateway acts as a multi-tenant network service connecting Trustgrid nodes in 100s or 1000s of unique customer environments with a fabric of connectivity that is managed through a single pane of glass.

Typically customers seek Trustgrid when they are migrating a legacy application to the cloud and realize that managing dozens or 100s of siloed VPNs does not scale efficiently.

Trustgrid gives these SaaS administrators and DevOps teams the ability to manage the networking component of the service in the same way they operate the other cloud services.

9B On-premise hardware supported by cloud services

Not all applications are destined for the cloud. Sometimes latency, customer preference, or deployment challenges mean software will continue to be deployed on a hardware appliance and shipped to a customer. However, even for these hardware dependent deployments, many still want to have some ancillary services running in the cloud for ease of maintenance and support.

The Trustgrid platform enables these applications to add cloud services without requiring any change to the application. Running on either a dedicated Trustgrid network appliance or deployed directly to the application's hardware appliance, Trustgrid can give an application provider global visibility and control of their remotely deployed appliances. These features can be used to remotely push new application versions to their appliances, monitor application performance, and even remotely deploy, support, and troubleshoot application components throughout their lifecycle.

Trustgrid's customers typically seek this type of a solution when they want to add new cloud-enabled services to applications that will continue to run on-premise or gain operational efficiencies in managing their fleets of remote appliances.

9C Cloud to edge architectures with Docker containers

Some applications may desire a more flexible architectural approach that does not limit delivery, features, or roadmap to a single premise. In these distributed or hybrid architectures, some application workloads will process at the edge and some will run in one or more public clouds, with data moving securely between all nodes.

For reasons of network latency, difficulty moving edge data, or lack of a need to ingest all remotely produced data to a centralized location, these distributed applications will need to run some workloads in IT environments they control and some behind firewalls they can't control. When this is desired, Trustgrid gives software providers the ability to seamlessly build, connect, and deliver applications across any environment (owned or 3rd party) with the same levels of control they experience in the cloud.

The Trustgrid platform enables this by supporting networking and edge computing features on the same virtual appliances. The Trustgrid nodes support proprietary scripts, ETL functions and Docker containers across any cloud or on-premise environment.

From within the platform, existing APIs can be secured with network protection and new APIs can be created from edge data sets. The Trustgrid management portal is used to manage scripts and containers, push and track versions, and provide all troubleshooting tools for remotely deployed software components.

Trustgrid's customers typically seek this type of a solution when they want to deploy code across any environment, but manage it as if it is a public cloud service. This allows their software development and DevOps teams to remove themselves from the provisioning of networking and management of edge computing capabilities, and focus on delivering software to hundreds or thousands of locations with minimal support hassles.

10 Conclusion

Delivering applications from the cloud or data center has become more complex than ever.

Today's distributed application architectures have evolved from the era where data and compute lived behind the same firewall, to architectures dependent upon remote access to customer-controlled data, microservice APIs, and on-premise deployments integrated to cloud services.

While these architectures can deliver improved performance, economies of scale, and enhanced customer experience, implementation of these architectures is dependent on several external factors that are not in the control of an application provider.

The Trustgrid platform recognizes the need for SaaS providers to control and support their application from the cloud to the edge.

Leveraging advanced networking, security and container based edge computing capabilities, Trustgrid allows SaaS applications to modernize their infrastructure without compromising their product or speed of innovation.

888.308.8995
info@trustgrid.io

trustgrid