eG Innovations

Achieving End-to-End Visibility: Best Practices for Monitoring Azure & AVD Environments

www.eginnovations.com

Executive Summary: Visibility is Mission-Critical

Remote and hybrid work models have fundamentally reshaped modern business operations.

With the ubiquity of internet access and a vast ecosystem of web-based tools and platforms, organizations have been able to scale back their operational costs, offer employees greater flexibility, and cast a wider net when it comes to recruitment.

Azure Virtual Desktop (AVD) has emerged as a critical enabler of this shift, allowing for flexible, scalable digital work environments and secure access to applications and data from virtually anywhere.

But AVD also presents unique challenges for IT and management teams. Native Azure tools leave critical visibility gaps that can make it difficult to detect issues and resolve UX or performance problems. At eG enterprises, we help organizations close these gaps with purpose-built end-to-end monitoring systems.

In this eBook, we'll explore proven strategies and best practices for gaining full visibility into your Azure and AVD environments. With these tips and tricks, you can identify issues faster, optimize resource usage, and deliver a seamless experience to every user, wherever and whenever they work.

Table of Contents

Chapter 01

Why Azure Observability is Critical for Modern IT

IT performance is a frontline business priority. And for enterprises that rely on Microsoft Azure, observability is paramount. To make sense of the data is to understand the system, and that understanding enables smooth sailing.

It's the linchpin to ensuring secure, seamless user experiences across today's complex IT landscape, where cloud–native, hybrid, and remote-first environments intersect and evolve at breakneck speed.

But what exactly do we mean by "observability?"

In simple terms, it refers to the ability to understand the internal state of a system based on the data it produces.

Your system's logs might be a good starting point, but they often fail to tell the whole story — you need a cohesive, contextual view that ties them together in real time.

Basic monitoring might alert you when something goes wrong. But a system built for observability allows your IT teams to delve deeper and uncover the "why" behind performance anomalies. In Azure environments, observability becomes even more critical, as infrastructure, applications, user sessions, and data often span multiple services, regions, and endpoints.

A modern Azure deployment might include a mix of:

- Virtual machines
- Containerized applications
- Databases
- API gateways
- Identity services
- AVD
- Third-party integrations

Each component is essential, and an issue in one area can affect performance elsewhere in the system.

For example, if an Azure SQL database experiences latency, AVD users may see slow load times or timeouts. With basic monitoring, you might get alerted to a high CPU warning on the SQL server. But those alerts can't tell you why it's happening or how other components are being affected. For IT teams, this makes for time-intensive diagnostics based on guesswork.

A system built with observability in mind avoids these pitfalls. Rather than isolated metrics, you get a bird's-eye view of your digital environment. You can trace the performance issue from the AVD session back to the SQL database, see latency spikes in real time, and understand how that's impacting downstream services. You might even see the root cause before the reports start pouring in.

As your IT environment grows more complex, observability features become essential to system maintenance: When your IT team has the at-a-glance insights they need to quickly resolve issues, your organization can stay agile and productive.

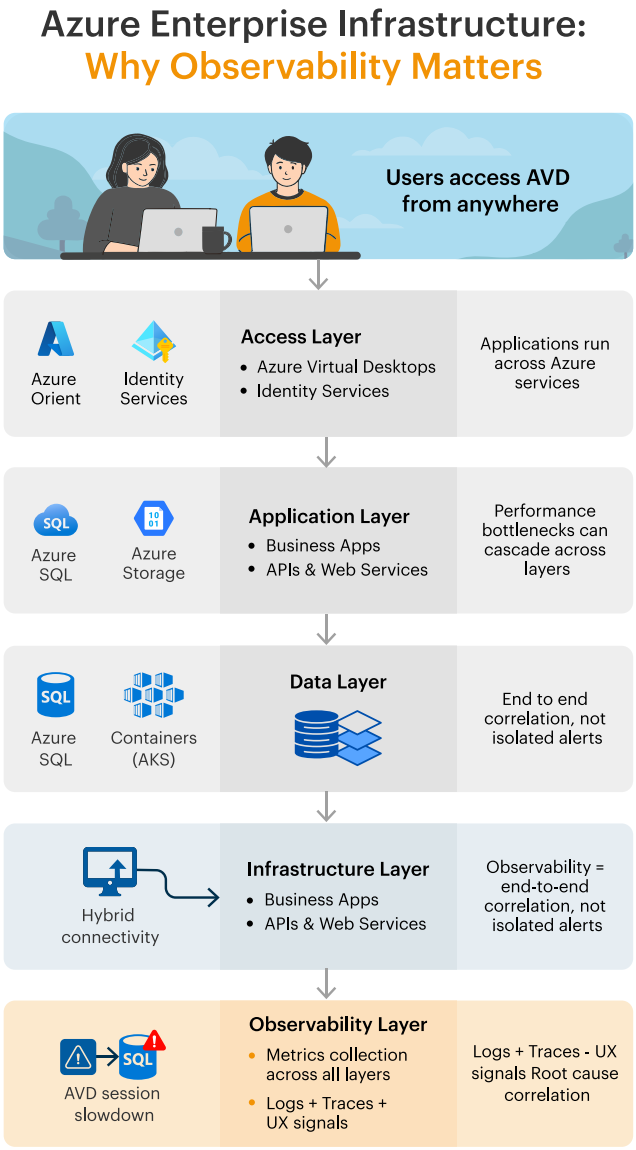

Figure 1: Caption to be updated

Chapter 02

The Strategic Role of Azure Virtual Desktop (AVD)

Azure Virtual Desktop is among the leading Desktop-as-a-Service (DaaS) solutions for modern enterprises. It enables secure, scalable, centralized management of virtual desktops and provides users with consistent access to apps and data whenever or wherever they're working.

AVD runs on Microsoft Azure, allowing IT teams to provide a digital workspace without shipping physical hardware or managing local installations. But while AVD simplifies delivery of remote work environments, it introduces new monitoring and management challenges that traditional tools can't handle.

Performance and UX (user experience) issues in AVD can be hard to pinpoint. Native Azure tools typically provide only surface-level data, meaning IT teams often struggle to get to the bottom of system disruptions.

Multiple layers of infrastructure, complex networking interdependencies, and roaming user profiles can create a fragile chain where issues in one area cascade into widespread performance problems, and it can take considerable time to comb through the data and identify the root cause.

Latency, logon issues, and profile misconfigurations account for the overwhelming majority of helpdesk tickets and are often the result of deeper infrastructure inefficiencies. This can eat up critical resources

within organizations — users are unable to work efficiently, and IT teams scramble to deal with the issue.

For example, a single misconfigured FSLogix profile can delay logins across an entire department, delaying project work and pulling IT away from other priorities. With proper observability in place, the problem might be identified and resolved early. For instance, an alert could flag abnormal profile load times for a specific host pool, prompting IT to investigate and apply a fix before users even notice.

Unfortunately, Azure native tools lack the deep, session-level visibility needed to mitigate such an issue. Without a more robust visibility solution, IT teams are often left playing whack-a-mole instead of focusing on long-term improvements. On the other hand, a system built with end-to-end visibility can turn fragmented signals into actionable insight across the full AVD stack.

Figure 2: Caption to be updated

Chapter 03

Common Performance Monitoring Challenges in Azure & AVD Environments

Monitoring Azure and AVD environments presents a unique set of challenges for enterprises. The complexity of modern IT environments and the lack of granularity in native monitoring tools make it challenging for IT teams to identify and address performance issues. Here's a closer look at the problems and solutions that define the day-to-day reality of managing these systems.

Hybrid Deployment Complexity

Enterprises typically run both on-premises and cloud infrastructure. Monitoring both often requires separate tools and dashboards, meaning critical insights are siloed in their respective platforms and hidden beneath layers of data.

This fragmentation prolongs the troubleshooting process and limits the ability of an IT team to make proactive performance and UX optimizations.

Limited Visibility Across Services

Azure environments are made up of many interdependent services, including VM, storage, networking, identity services, and third-party integrations. Native tools can only monitor these components in isolation.

That makes it hard for IT teams to get a system-wide view of performance. As a result, issues often go undetected until they start causing problems for users.

Performance Bottlenecks

Underprovisioned VMs, storage limitations, and overloaded host pools are among the most common culprits behind bottlenecks in AVD environments, which can frustrate users and cause headaches for IT teams. It's not always immediately clear where the slowdown is coming from, especially when surface-level metrics suggest everything is functioning normally.

Poorly Tuned Alert Thresholds

IT teams are often bombarded with alerts that aren't properly prioritized or tied to relevant business impact. For example, a minor CPY spike on a single VM might trigger the same level of alert as a login failure across an entire host pool. Without well-calibrated alert thresholds, it's harder for IT teams to sift through the noise and respond to what actually matters.

Root Cause Analysis Difficulties

Troubleshooting across multiple layers can be time-consuming, often involving multiple tools, dashboards, and teams. Minor outages may require a protracted process of elimination, and major outages requiring a more in-depth RCA can stall operations while teams scramble to piece together fragmented data and timelines.

What's at risk?

Visibility limitations often cause a ripple effect of disruption that extends well beyond the IT department:

- **Lost productivity:** Laggy sessions, login delays, and slow applications can disrupt day-to-day work and frustrate users. Oftentimes, they may "accept" poor performance as part of the normal remote experience, meaning issues go unreported.
- **SLA breaches:** Unresolved performance issues can violate internal or external service level agreements, leading to missed contractual obligations and even penalties. This can potentially damage credibility with customers and stakeholders alike.
- **High MTTR:** Incidents take longer to diagnose and resolve without clear visibility. This strains IT resources, increasing operational costs and limiting the capacity for IT teams to focus on maintenance and optimization.

Chapter 04

Best Practices for Monitoring Azure & AVD

Effective performance monitoring of Azure requires a thoughtful strategy that accounts for the complexity of your IT infrastructure and user expectations. Use the checklist below to guide your monitoring setup and ensure end-to-end visibility across your environment.

Scale collectors appropriately

To capture telemetry from both Azure and on-premises environments, you'll need to consider a few different factors when scaling your monitoring collectors:

- **Session volume:** The number of concurrent user sessions directly affects the data collection load, especially in AVD. Collectors must be scaled appropriately to ensure they can handle the increased volume without delays or data loss. It's advisable to establish a baseline (e.g., one collector per X concurrent sessions) and make accommodations for projected peaks and spikes.
- **Workload type:** Different workloads produce different types and volumes of telemetry, and you'll

need to tune your collectors accordingly. You can use workload profiling to, for example, dedicate one collector to a resource-heavy SQL server and group lower-traffic systems under a shared collector. Allocating resources in this way keeps your system balanced and prevents bottlenecks.

- **Geographic distribution:** If your system spans multiple geographic locations, it's best to deploy collectors close to the source of the data. Rather than relying on a central collector, deploy at least one per active Azure region to localize data insights. This helps reduce latency, improve accuracy, and ensure region-specific issues are detected quickly and in context.
- **Data retention policies:** It's important to consider how long you store your telemetry and at what resolution. Higher-frequency data or more extended retention periods require more throughput, so it's wise to balance granularity and history based on use case. For example, you might store per-minute metrics for one week to aid troubleshooting, and hourly data for 90 days to support trend analysis and planning.

Correlate metrics across time, layers, and users

Tracking individual metrics in isolation can highlight symptoms but rarely reveals the full picture. To truly understand system health, you'll need to correlate data across infrastructure layers, user sessions, and timeframes. This allows for faster, more accurate diagnostics and a clearer understanding of how issues can affect users.

For example, a latency spike in an AVD session may not seem significant until you see it lines up with a CPU bottleneck on a specific VM. This kind of correlation is crucial for helping IT teams stay on top of performance issues and prevent situations where they might end up flooded with helpdesk tickets.

Separate infrastructure alerts from user experience signals

Not all alerts amount to a five-alarm fire. To prioritize effectively, you'll want to create a triaged alert system, differentiating between infrastructure-level alerts (e.g., high CPU, disk latency, network throughput) from user experience signals like slow logons or session lag.

Tiering alerts in this way allows your team to respond with an appropriate level of urgency — system-wide or user-facing issues can be handled immediately. In contrast, less critical alerts can be monitored and resolved as part of your routine maintenance workflow.

Track the KPIs that matter most

To maintain system performance and resolve issues before they impact users, you'll want to keep the most relevant KPIs front and center. We recommend having the following metrics actively monitored and visualized in your main dashboard or alerting system:

- **CPU ready time:** If this consistently exceeds 5–10%, your VMs are likely underprovisioned or overcommitted. Look for patterns during peak hours and consider resizing or spreading workloads.
- **Logon duration:** A good target is under 30 seconds. Spikes may signal FSLogix delays, GPO bloat, or authentication bottlenecks. Break this down by stage (profile load, group policy, desktop load) for more precise insights.
- **Session latency:** Anything above 100ms may be noticeable to users, and over 250ms typically triggers complaints. Use latency heatmaps by region or host pool to pinpoint where degradation is occurring.
- **App load time:** Consistently slow load times may signal storage performance issues or profile bloat. Compare across users and session hosts to distinguish user-side vs. infrastructure issues.
- **Profile load time:** If profile load exceeds 20 seconds, investigate FSLogix container performance, profile size, and file share throughput. Look for growing trends, not just outliers.

- **Connection failure rates:** Even a slight uptick here can be meaningful. Correlate with host availability, broker errors, or conditional access policies to uncover hidden access issues.

Set up purpose-built monitoring to streamline visibility

Generic monitoring tools often miss the nuances of Azure environments, but third-party solutions can bridge the gap. eG Enterprise offers a purpose-built monitoring solution with out-of–the-box support for:

- AVD session hosts
- FSLogix profiles
- Connection brokers
- EntraID

Getting started is simple. Deploy lightweight collectors in your Azure and on-prem environments, auto-discover components using built-in templates, and begin monitoring with minimal manual configuration. From there, you can customize your dashboards, set intelligent alert thresholds, and start tracking end-to-end performance from a single pane.

Want to see it in action?

To explore how eG Enterprise can help bring full-stack visibility to your Azure systems.

Request a demo

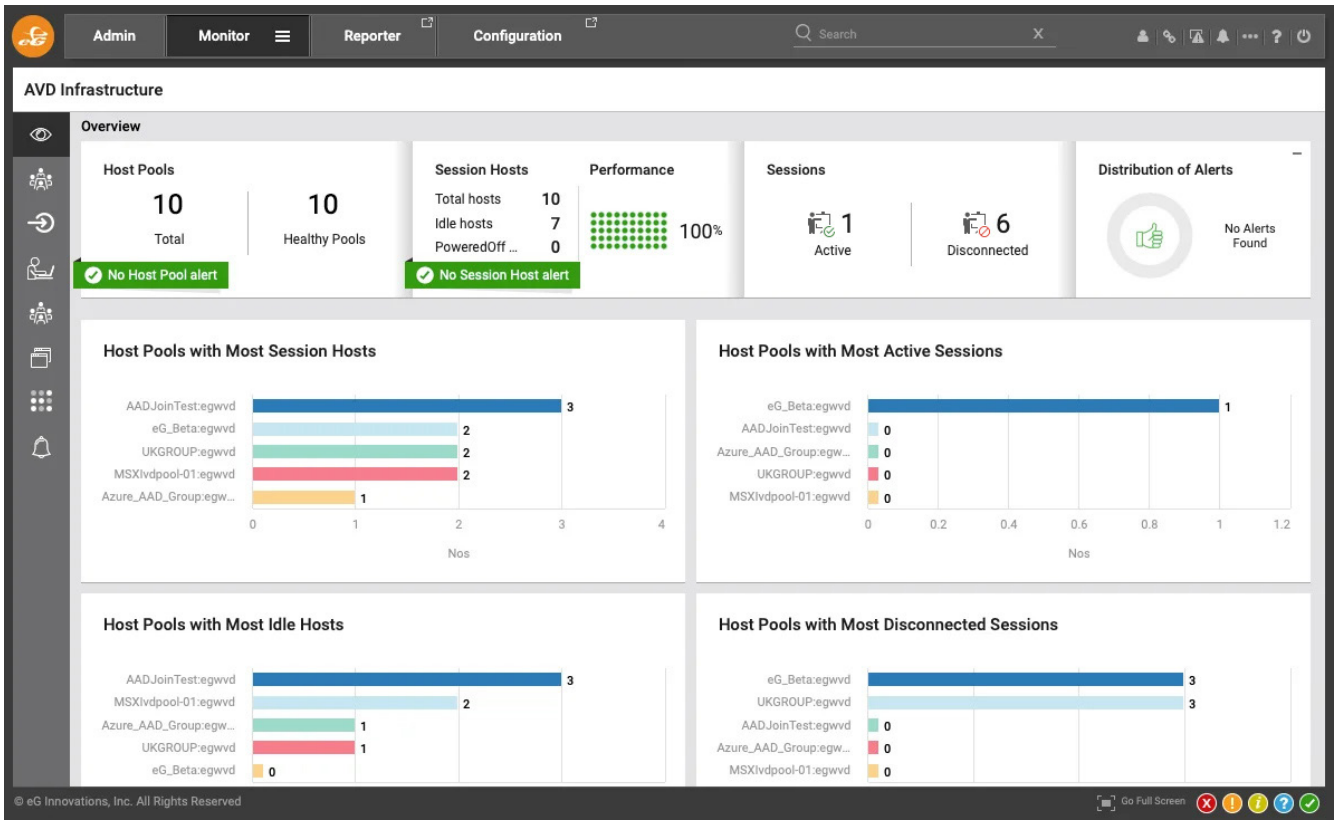

Figure 3: Caption to be updated.

Chapter 05

Where Azure Native Tools Fall Short

Microsoft's native tools for monitoring Azure environments — Azure Monitor, Azure Insights, and Log Analytics — provide a solid foundation for collecting telemetry, setting alerts, and performing custom queries. But for enterprise-level organizations, these tools often fall short. Here's why:

Siloed Metrics

Azure's native tools tend to separate application, infrastructure, and UX data into different dashboards. VM performance may live in one view, application health in another, and user behavior in yet another.

This makes it difficult to understand how issues in one place affect performance in another. It can make for a needlessly long diagnostic process and a limited capacity for optimization.

Steep Learning Curve

Effectively using native Azure tools often means learning Kusto Query Language (KQL) and understanding the nuances of multiple Azure dashboards. This can be especially challenging when managing large-scale, multi-layered environments.

Simply put, many IT teams don't have this specialized knowledge or the budget to hire Azure specialists.

Lack of Root Cause Correlation

Azure Monitor and Log Analytics can raise alerts but don't connect the dots between symptoms and causes. With no built-in end-to-end tracing, pinpointing the "why" behind performance problems often requires manual investigation across multiple data sources. This not only slows down resolution time but also increases the risk of misdiagnosis.

Limited User Experience Insights

Native Azure tools are designed primarily for infrastructure monitoring, not for tracking how real users interact with the systems they use. Metrics like logon duration, session latency, or application launch times are often either buried or not available at all. That limits the ability of IT teams to proactively address performance issues or validate that services are delivering as expected.

Fragmented Alerting and Dashboards

Alerts, logs, and visualizations are often managed in different places, which impedes the ability of IT teams to respond quickly and makes holistic monitoring a challenge. Without a centralized view, teams must jump between tools to piece together context. This slows things down, increases the chance of missed signals, and makes coordinated incident response more difficult.

Complex systems demand a more robust solution that goes beyond isolated alerts and fragmented dashboards. They need an integrated observability approach built for modern, distributed environments — think end-to-end visibility and proactive root cause analysis.

Chapter 06

eG Enterprise for Unified Azure & AVD Observability

As enterprise IT environments grow more distributed and performance-sensitive, eG Enterprise offers a comprehensive observability platform engineered for today's Azure ecosystems.

Monitoring Azure Workloads

Modern Azure environments span a wide array of services, each with its own telemetry, performance benchmarks, and failure modes. eG Enterprise brings them all into a unified view so that IT teams have the clarity they need to work efficiently.

- **Track performance across Azure services:** eG Enterprise monitors key Azure components, including VMs, web apps, SQL databases, containers, storage accounts, and more. No matter how complex or distributed your deployment, eG provides real-time telemetry and intelligent baselining for each service.
- **Out-of-the-box integrations for fast deployment:** Built-in support for a full range of Azure services — Azure App Service, Azure SQL, AKS, and Azure Storage — means teams can hit the ground running. Get pre-configured dashboards, KPIs, and alert thresholds without the need for manual setup.
- **Flexible deployment options:** eG Enterprise supports both agent-based and agentless monitoring, meaning you can choose the approach that best aligns with your security and operational policies. Automated discovery makes it easy to onboard new services and scale coverage as your IT footprint grows.

Proactive Azure Service Monitoring

eG Enterprise goes beyond traditional monitoring by using machine learning and context-aware alerting to keep teams focused on what matters.

- **Auto-baselining for dynamic environments:** Performance thresholds aren't static. eG Enterprise automatically learns your system's normal behavior over time and adjusts baselines dynamically. This helps reduce false positives and allows for a smarter alerting framework that surfaces only genuine anomalies.
- **Catch issues before users notice:** Predictive analytics and trend-based alerting can detect potential service disruptions well before they degrade performance. This allows you to fix problems proactively without relying on help desk tickets to signal that something's gone wrong.
- **Cut through alert noise with correlated insights:** Unlike native Azure tools that fire off isolated alerts, eG Enterprise correlates events across infrastructure layers. For example, if a slowdown is tied to both a storage latency spike and high CPU usage on a backend VM, eG Enterprise flags the relationship so your team can tackle the problem efficiently.

Specialized AVD Monitoring

Virtual desktop environments come with their own set of performance risks, and users tend to notice issues right away. eG Enterprise is purpose-built to address the unique demands of AVD.

- **Deep session-level visibility:** Track real-time and historical metrics for each user session, including logon duration, profile load times, session latency, and application launch delays. eG's dashboards help your team pinpoint UX problems fast without digging through multiple tools.
- **Multi-layer drill-downs:** Diagnose issues across every layer of the AVD stack: from the user session and VM to the broker, gateway, FSLogix profile, and network path. With eG, you can zoom in and out easily — a single click can take you from a high-level performance dashboard down to the root cause.
- **Data-driven optimization:** By analyzing session trends and infrastructure usage, your team can make informed decisions about host pool sizing, address recurring performance issues, and reduce login times. These continuous improvements lead to a more efficient environment and significantly enhance the experience for end users.

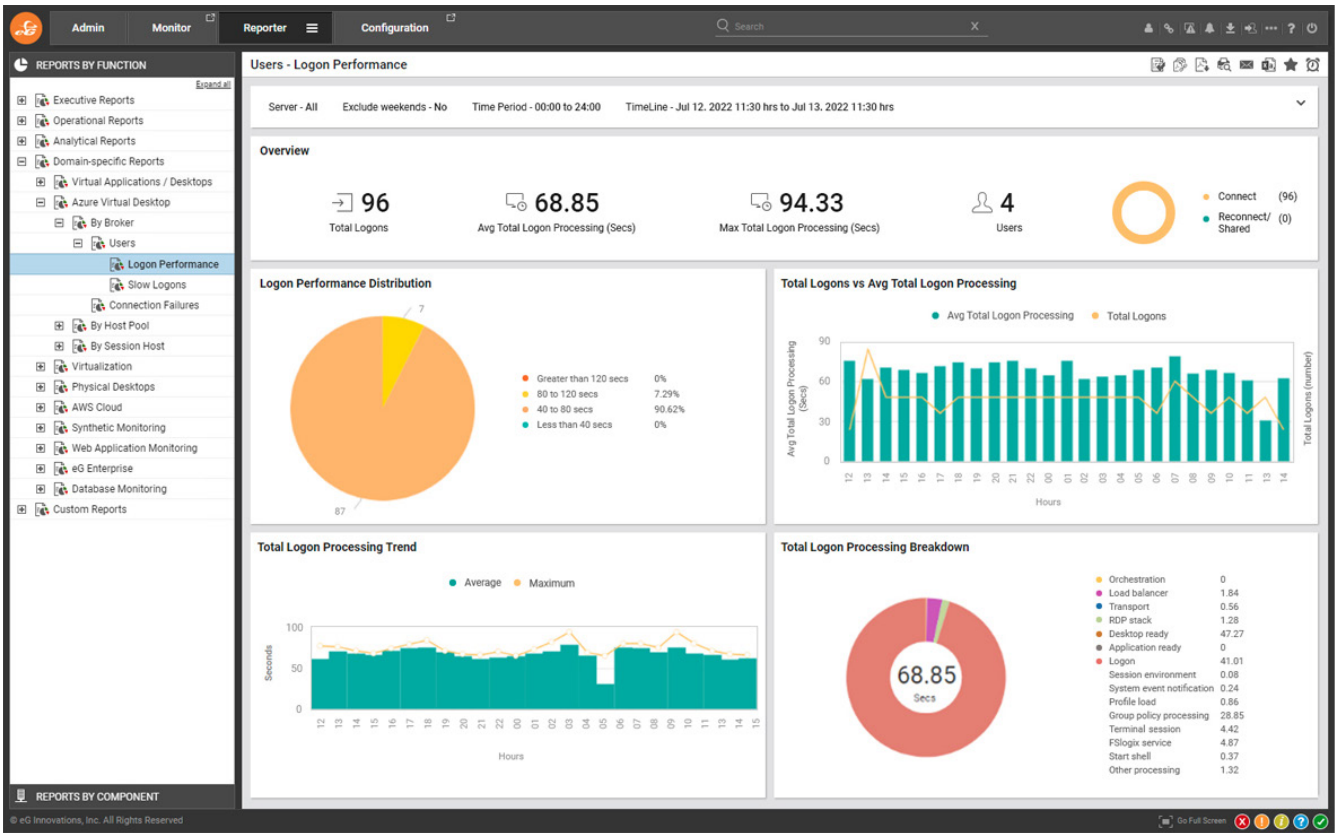

Figure 4: Caption to be updated.

Conclusion

Take Control of Your Azure & AVD Monitoring

Your IT environment demands more than basic metrics and disconnected dashboards. With eG Enterprise, you get unified observability across Azure workloads and AVD infrastructure, empowering you to diagnose root causes faster and prevent downtime.

From session-level insights and auto-baselining to intelligent alerting and drilldowns across layers, eG Enterprise delivers the end-to-end visibility your IT team needs to keep your systems running and your workforce connected.

See eG Enterprise in Action

Schedule a Demo

About eG Innovations

eG Innovations is dedicated to helping businesses across the globe transform IT service delivery into a competitive advantage and a center for productivity, growth and profit. Many of the world's largest businesses use eG Enterprise to enhance IT service performance, increase operational efficiency, ensure IT effectiveness and deliver on the ROI promise of transformational IT investments across physical, virtual and cloud environments.

To learn more visit www.eginnovations.com

eG Innovations

+1 866 526 6700

info@eginnovations.com

www.eginnovations.com