White Paper

Mobile Network Performance Visibility

The journey from 4G to 5G standalone in 5 steps

Introduction

The journey from 4G to 5G mobile networks involves more than just the additional spectrum, higher bandwidth, and better connectivity. At its heart, it is a cloud migration project for mobile operators, where edge cloud performance plays a key role for 5G services. This is similar to the cloud transformation journey in many enterprises and organizations.

All core pieces of mobile networks, whether network control functions, authentication, billing, roaming handovers, or Internet connectivity, will be deployed as cloud-native software designed to automatically scale and be resilient to infrastructure failures.

For mobile operators, the main business drivers behind these changes and the move from 4G Evolved Packet Core (EPC) to a 5G core network are:

- New types of services and revenue streams (e.g., fixed wireless broadband, massive-scale IoT, low-latency applications that leverage edge computing)
- Faster time to market for new services
- Decreased vendor lock-in by deploying generic compute infrastructure
- Reduced operating expenditure with an artificial intelligence and machine learning-assisted fully-automated network

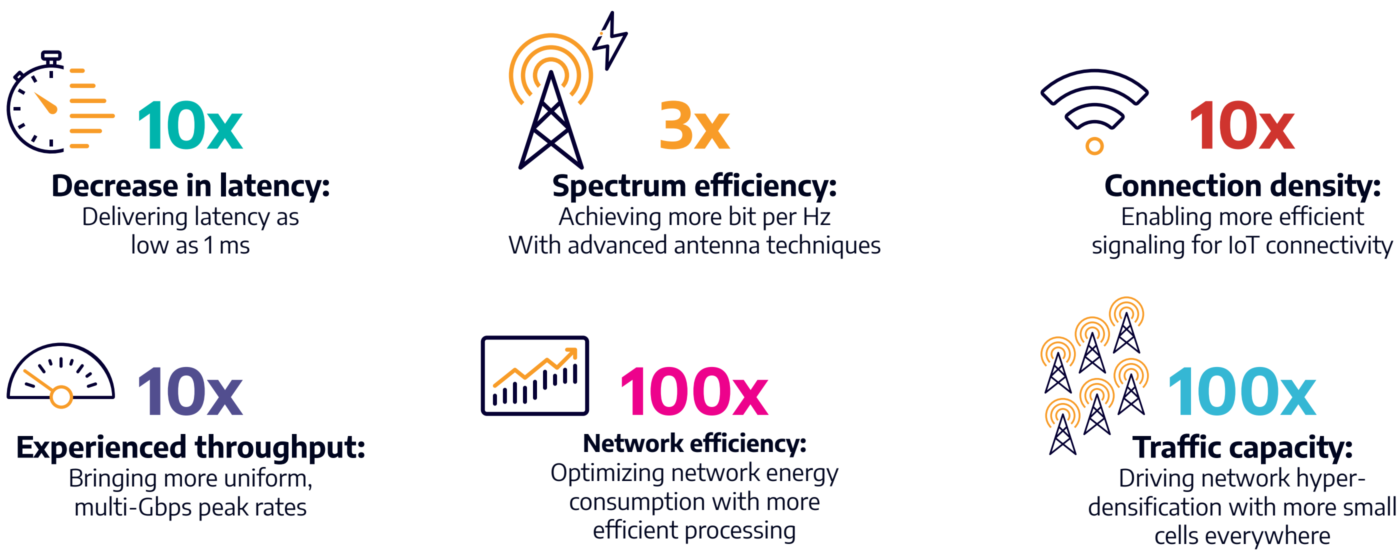

Figure 1: The technical promises of 5G are underpinned by performance

Naturally, 4G operators can't switch to a fully-blown 5G network overnight, but carefully selected steps will minimize the impact on existing revenue and operations while introducing 5G flexibility, service enablement, and automation.

Much like enterprises have employed DevOps models for continuous improvement of their services, many mobile operators choose this path for their 5G journey. A vital component of successful DevOps initiatives is testing. **Testing, and then testing again, deploying, and then testing.**

The entire service lifecycle, including service assurance, has to be as automated and dynamic as possible, allowing for continuous quality of service (QoS) and quality of experience (QoE). For example, adapting network resources to changing network conditions in order to maintain high QoS and QoE.

Performance is the key to unlocking much of the value that 5G networks will bring to mobile operators. Adopting solutions capable of automating detection, analysis, and action upon performance deviations is imperative in building the self-healing 5G service overlay.

From 4G to 5G in steps, assuring performance along the way

Deploying and maintaining 5G mobile networks are a massive undertaking that involve transforming the network, IT systems, and organization skills based on a phased strategic plan rather than a big bang approach. An operator may take a different path based on their market and customers and decide to skip some steps or prioritize others, but the fundamental 5G building blocks remain the same.

In the next sections, we discuss the transformative steps to 5G and the critical performance questions to answer to deliver exceptional customer experience along the way.

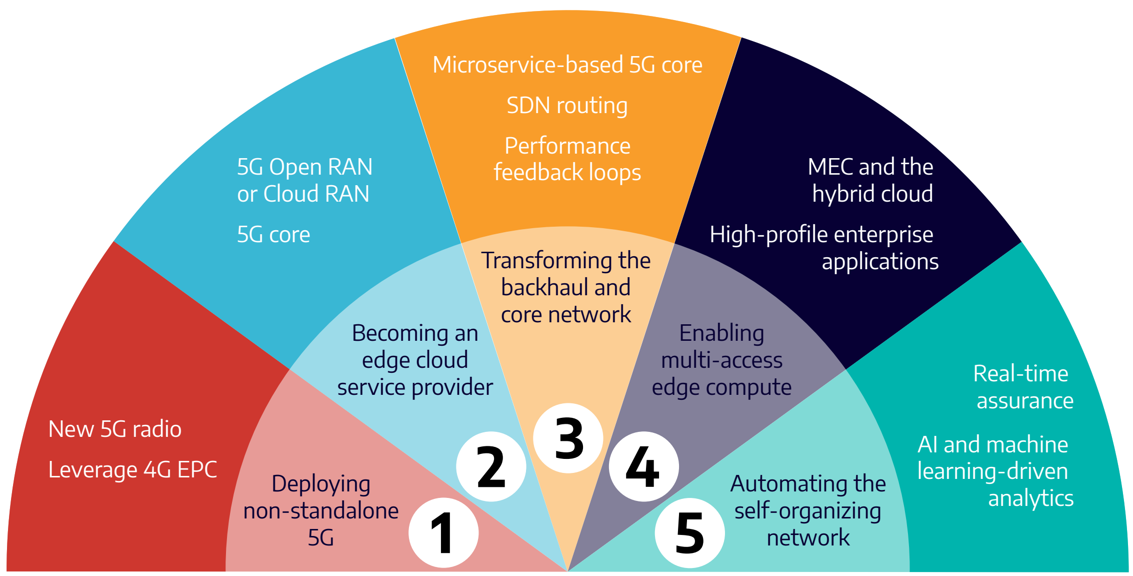

Figure 2: The 5 steps towards standalone 5G, maintaining assurance and customer experience along the journey

Step 1: Deploying non-standalone 5G

Leveraging existing backhaul network capacity and adding 5G New Radio (NR) is the natural first step on the 5G journey.

Many operators are already in trials or in the early stages of deploying non-standalone (NSA) 5G. Benefits include winning the marketing race and being one of the first 5G operators on the map, but also the introduction of new frequency bands increases bandwidth available for customers. Bandwidth and spectral efficiency are what NSA 5G is essentially about.

Wireless broadband for residential and business customers can now be provided with Gigabit speeds (assuming the existing backhaul can cope with the added data). This positions 5G fixed wireless access as a serious competitor to broadband, especially in areas with older traditional copper (DSL) or coax cable services.

The addition of 5G NR also has benefits for the Internet of Things (IoT) market. The new radio allows far more connections per square kilometer (one million vs. two thousand with 4G),[1] a key enabler for cheap and low-bitrate (Narrowband IoT) connections. New 5G base stations can also be combined with small cells or active antenna solutions to improve indoor coverage.

Overall, this first foray into 5G adds considerable new traffic to the network that needs to perform just as well or even better than the existing 4G network. Customers certainly expect to turn on their brand new 5G phone and get a better experience than their old LTE handset.

Mobile operators need to test performance before deploying 5G NSA to validate continuously that data paths and data processing services are working at a high level of efficiency and performance. If you are adding more bandwidth to your backhaul infrastructure to cope with 5G NSA, you need to add it at the right time and place to ensure it is providing the desired performance improvements.

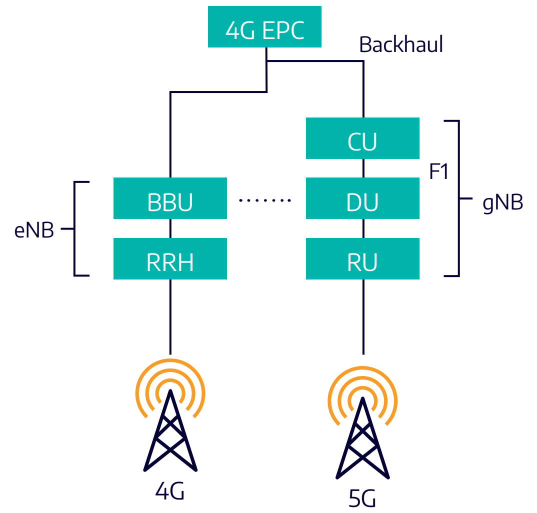

Figure 2: Non-standalone 5G RAN leverages the existing 4G backhaul and 4G Evolved Packet Core (EPC)

Top 5G NSA performance questions to answer:

- Do you have accurate performance measurement to the eNB/gNB sites (backhaul/midhaul or F1?) before the site is turned up?
- Do you have a continuous accurate measurement of Layer 2-4 connectivity from the core to the eNB/gNB? Including delay, packet loss, and throughput?

[1]https://www.vodafone.com/what-we-do/technology/5g

Step 2: Becoming an edge cloud service provider

The introduction of 5G Open RAN (O-RAN) or Cloud RAN (C-RAN) and the 5G Core Network (CN) will require new compute platforms in the core, but also at the edge of the mobile network.

With edge computing and new 5G network functions, mobile operators can handle data traffic much more flexibly and efficiently than before. Specific application and user or service data flows can be separated, sorted into slices, and managed separately from other traffic types. The ability to filter out generic internet traffic right at the base station and peering the traffic directly to the internet offers tremendous benefits to mobile operators and customers. This results in the backhaul network not getting loaded unnecessarily so customers see faster response times on their internet services.

This intelligent traffic separation and edge computing ability also offers a massive opportunity for mobile operators to deploy edge data centers close to end users.

For the first time in history, imagine enterprises being able to upload and run programs or partial workloads inside the mobile network, just a millisecond away from the user. End users would be able to instantly upgrade their quality of service or select the internet peering points that give the best gaming or augmented reality experience right at that moment.

This may not happen right at the early stages of deploying 5G, but the enablers of edge computing and application separation are the starting point.

5G edge clouds will be able to hand-off specific user or application traffic to one of the existing web-scale cloud providers: Amazon Web Services (AWS), Google, Alibaba, Tencent, or Microsoft. These cloud providers are introducing new edge services to support the market's development, including AWS Wavelength, Google's Global Mobile Edge Cloud and Anthos, and Microsoft's Azure Edge Zones.

However, a mobile operator's business model is more interesting if it can directly provide and sell at least part of this service. Many 5G operators are starting to trial and deploy mobile edge platforms and portals to offer localized compute infrastructure. The new proximity of the edge service will enhance the QoE for end users for applications like augmented and virtual reality, cloud gaming, and mobile advertising (read more on this in Step 4: Enabling multi-access edge compute).

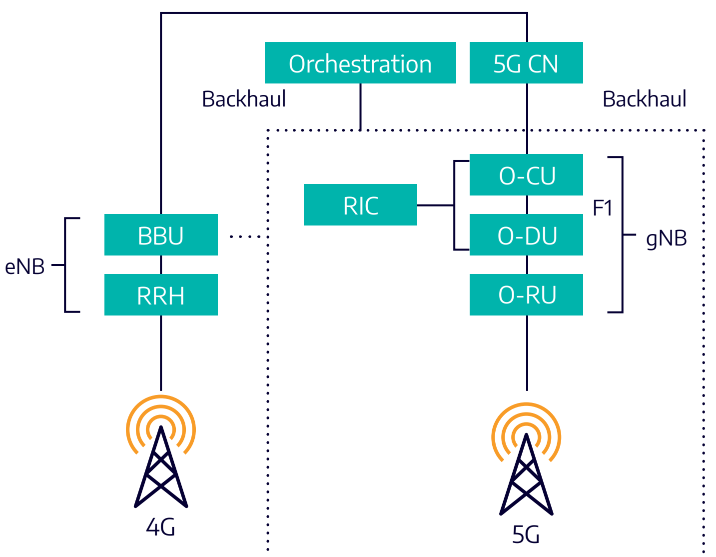

Figure 3: The move to 5G Standalone (SA) is designed to leverage the agility of an open software-driven network, plus cloud-native O-RAN and 5G CN

Top 5G SA performance questions to answer:

- Do you have accurate performance measurements to the ng-eNB/gNB sites (backhaul/midhaul or F1?) before the site is turned up?
- Do you have a continuous accurate measurement of Layer 2-4 connectivity from the core to the ng-eNB/gNB? In a mesh between gNB sites? Including delay, packet loss, and throughput?
- Can you measure the network performance to and from the cloud (MEC) infrastructure supporting the O-RAN/5G CN components?
- Can you measure the application performance of the O-RAN and 5G CN components? HTTPS? Event bus? Can you measure the performance of supporting infrastructure (e.g., DNS, etc.)? Inside of Docker/Kubernetes clusters?

Step 3: Transforming the backhaul and core network

5G allows for more flexible resource usage where major radio and network functions can be centralized and implemented using off the shelf hardware, yet distributed closer to the user as needed. Deploying Cloud-RAN or Open-RAN, both centrally and in a distributed manner, will be the key to success. It is important to assure that the connectivity is top-notch and ideally feeding network performance KPIs to the software-defined networking (SDN) controllers of the overlay routing and path selection systems.

A working 5G core network puts new requirements on the transport network infrastructure. Routing and switching needs to be programmable and efficient with feedback loops to detect and secure performance consistency and integrity.

A proper implementation of both core and edge cloud environments will enable successful 5G cloud-native packet core operations. This includes the ability to implement multi-access edge computing where third parties can deploy and orchestrate workloads inside the mobile network and benefit from the scale and agility of a microservice-based 5G core network.

However, new tools will be required to understand the integrated performance of the edge cloud, access and core networks, which will be critical to offering differentiated 5G services. Visibility into the performance of 5G O-RAN, 5G core network and third-party applications and supporting services (e.g. DNS, etc.) will be critical.

Acronym Key

Let us help you navigate the storm of 5G acronyms.

- **3GPP** Third Generation Partnership Project
- **BBU** Baseband Unit
- **CN** Core Network
- **C-RAN** Cloud Radio Access Network
- **CU** Centralized Unit
- **DNS** Domain Network Server
- **DSL** Digital Subscriber Line
- **DU** Distributed Unit
- **eNB** Evolved Node B
- **EPC** Evolved Packet Core
- **F1** interface between the CU and DU
- **gNB** Next Generation Node B
- **IoT** Internet of Things
- **LTE** Long Term Evolution
- **NWDAF** Network Data Analytics Function
- **NSSF** Network Slice Selection Function
- **NR** New Radio
- **NSA** Non Standalone
- **O-CU O-RAN** Central Unit
- **O-DU O-RAN** Distributed Unit
- **O-RU O-RAN** Radio Unit
- **O-RAN** Open Radio Access Network
- **PCF** Policy and Charging Function
- **QoE** Quality of Experience
- **QoS** Quality of Service
- **RIC** Near Realtime Intelligent RAN Controller
- **RRH** Remote Radio Head
- **RU** Radio Unit
- **SA** Standalone

Step 4: Enabling multi-access edge compute

Building on the early and more static mobile edge compute deployments, multi-access edge compute (MEC) is a true cloud-native containerized execution environment for a diverse set of applications.

High-profile applications include support functions for autonomous vehicles, augmented reality, factory automation, and new industrial IoT services — applications that require the lowest possible latency and highest possible reliability. With the right business model, the MEC framework can be a significant new revenue stream for mobile operators, but it is also challenging to assure services.

In order to properly assess the performance of a service that is fully automated and instantiated on-the-fly, the test and assurance function needs to be included with the service definition and orchestrated together with the service.

Many operators will also end up operating hybrid cloud MEC environments, combined with resources from a web-scale cloud provider (see Step 2: Becoming an edge cloud service provider). Assuring this hybrid cloud environment also presents challenges for gaining performance visibility into the hybrid cloud with workloads potentially running in many different locations.

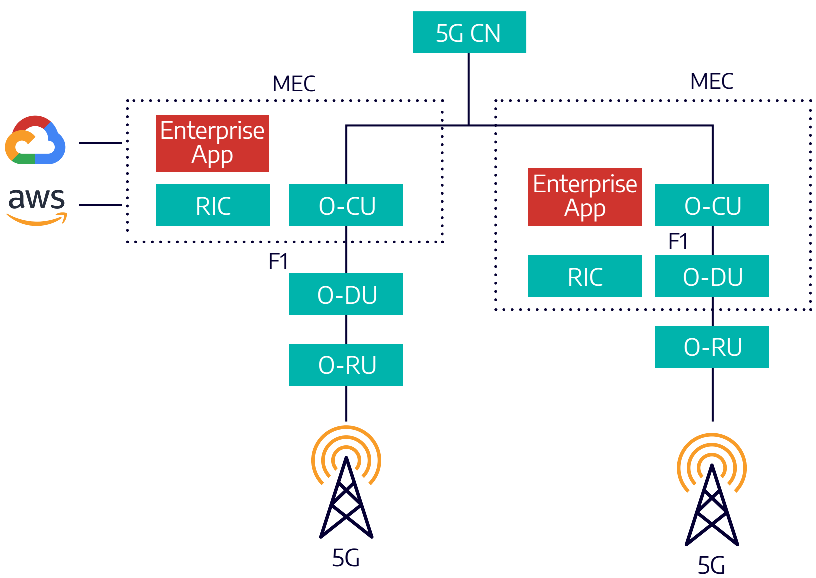

Figure 4: MEC leverages cloud infrastructure to host enterprise applications – application performance is foundational

Top MEC performance questions to answer:

- Do you have continuous accurate measurement of Layer 2-4 connectivity from the core to the MEC? Between MEC sites? Between MEC and public cloud providers? Including delay, packet loss, throughput?
- Do you have continuous accurate measurement of Layer 2-4 performance on a per-slice basis?
- Can you measure the application performance of the enterprise applications? HTTPS? Event bus?
- Can you measure the performance of supporting infrastructure (e.g. DNS, etc.)? Inside of Docker/ Kubernetes clusters?

Step 5: Automating the self-organizing network

The service-based architecture of a 5G network ultimately leads to the possibility of a completely self-managed network infrastructure. This is where intent-based configuration, using machine learning and AI algorithms, can automatically reconfigure network resources and automate performance testing to ensure the quality of customer experience is maintained. Indeed, the complexity of the 5G network requires that increasing levels of automation are put in place to gain operational efficiencies and ensure 5G can be profitable.

This means that all of the assurance functions for backhaul, midhaul, multi-access edge computing, and hybrid cloud need to be able to provide actionable insight in real time. This real-time insight can be leveraged by orchestration and automation functions for 5G infrastructure, and it can also be leveraged by the 5G core network.

The standards body 3GPP has created the Network Data Analytics Function (NWDAF) as part of the 5G core network definition. The NWDAF can provide real-time insight to other 5G core network elements (e.g. Network Slice Selection Function (NSSF), Policy and Charging Function (PCF), etc.) ensuring intelligent and automated decision making in the 5G core network.

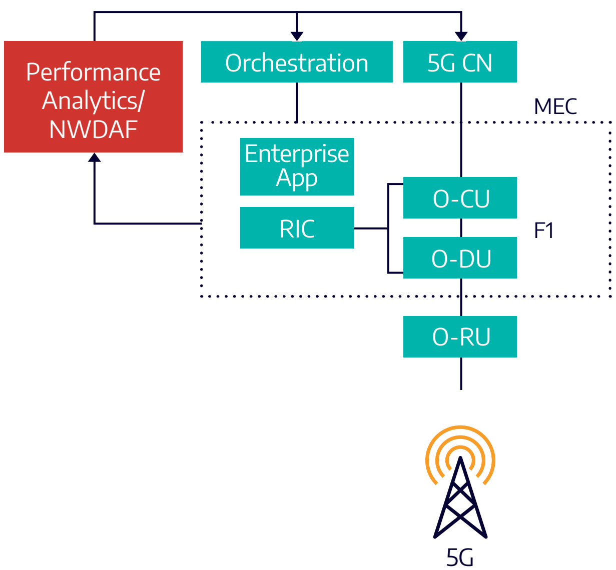

Figure 5: Self-organizing networks leverage real-time performance analytics to do intelligent orchestration

Top self-organizing network performance questions to answer:

- Can you collect performance data from multiple sources; from RAN to gNB to backhaul? From MEC and enterprise applications?
- Can you analyze and provide insight on this data in real time? (Very different from a “data lake”)
- Do you have open APIs and an event bus to provide real-time performance data to your orchestration layer? To the 5G CN elements?

Summary

For mobile operators, 5G is about more than just speed and spectral efficiency in the RAN. It's an opportunity to move to a cloud-native architecture throughout the mobile network (O-RAN and the core network) and to leverage that cloud architecture to offer unprecedented access to the mobile network for enterprises. 5G performance is going to be key, especially for the target market that will drive revenue: enterprise applications.

However, this new 5G network introduces more complexity and requires a very different performance monitoring proposition than in 4G. Besides the network underlay, application performance and real-time automation will all be needed to assure the end user experience.

The only way to manage the performance and complexity is to ensure end-to-end performance visibility. This means visibility from the RAN to the edge and core network and everywhere in between, including the transport underlay and the service overlay. This performance data has to be made available to the rest of the network for real-time decision making. Performance visibility will be key in unlocking both the financial and operational benefits of 5G network transformation.

Accedian Skylight for 5G

Accedian Skylight is ready to help you with 5G performance monitoring today and in the future; from network underlay to cloud and everywhere in between.

Accedian Skylight is a cloud-native performance management solution that provides a complete 5G network visibility of both the underlay and service overlay as an integrated solution. Skylight sensors can be deployed in the core, midhaul and backhaul networks, as well as between MEC and public cloud locations, to provide full mesh underlay performance visibility.

Skylight sensors can also be deployed in the MEC and cloud locations in order to provide full visibility into service performance for the O-RAN, 5G CN and enterprise applications. Skylight performance analytics can analyze performance data in real time and has open APIs for integration.

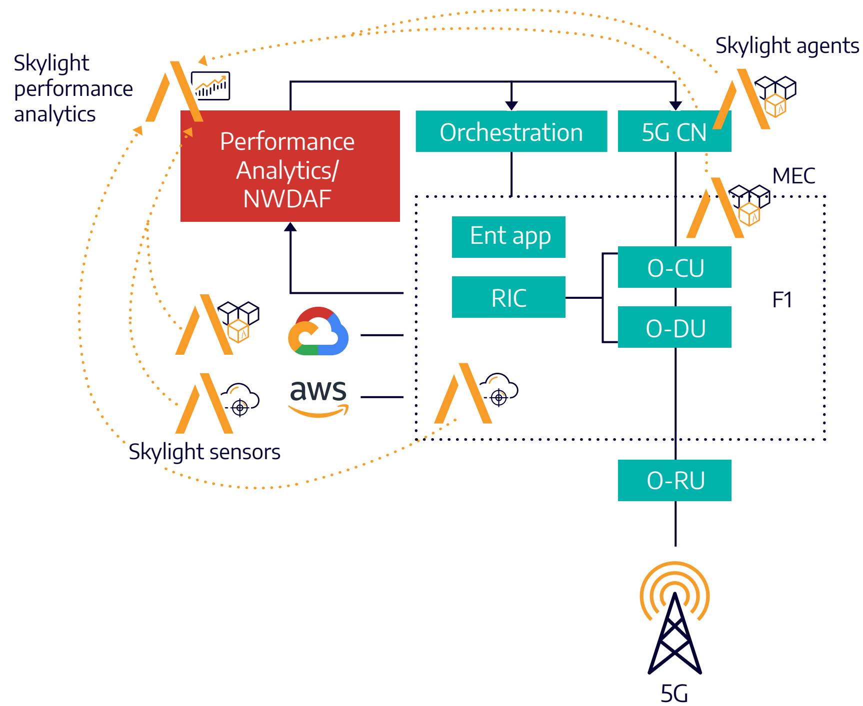

Figure 6: The cloud-native Skylight is a performance management solution that provides a complete 5G network visibility of both the underlay and service overlay as an integrated solution.

About Accedian

Accedian is the leader in performance analytics and end user experience solutions, dedicated to providing our customers with the ability to assure their digital infrastructure, while helping them to unlock the full productivity of their users.

Learn more at accedian.com

Accedian | 2351 Blvd. Alfred Nobel, N-410 | Saint-Laurent, QC H4S 2A9 | 1 866-685-8181 | **accedian.com**