COMMSCOPE®

Data Sovereignty: A Challenge for Any Data Center Anywhere

In the fable "The Tortoise and the Hare," a slow yet persistent tortoise challenges a hare to a race. The hare was far faster but also way overconfident. The slow, persistent tortoise kept moving along the path and, as we all know, won the race. A similar race is playing out worldwide as national governments slowly but surely progress toward data sovereignty objectives. Until recently, slow-moving bureaucratic institutions struggled to write legislation that kept pace with more nimble private entities that were quickly expanding their digital services across international boundaries. That has begun to change.

When the 27-member European Union (EU) adopted the General Data Protection Regulation (GDPR) in 2018, the new law initiated a wave of similar legislations that helped jump-start global efforts to ensure the sovereignty of personal data. GDPR defines personal data as any information related to a person that can directly or indirectly identify them. This includes name, location data, physical attributes, health information, ID numbers, and any other economic, cultural or social identity factors. It even includes online identifiers such as an IP address.

The introduction of GDPR provided a legal framework that enabled federal governments to, for the first time, levy financial penalties on corporations that had previously operated outside national jurisdictions. The potential severity of the penalties was a wake-up call for organizations to take the issue of data sovereignty seriously. Simultaneously, legislators in other countries and local governments started considering how they could use similar frameworks to protect their citizens and industries. However, national efforts such as these often come with unintended consequences.

Expanding Global Legislation

While large, multinational privacy regulations like GDPR make the headlines, countless smaller, regional laws receive less attention but create key obstacles for multinational businesses looking to expand. Adding to the complexity, the definition of "data sovereignty"—and how it applies to individuals, business entities and legal transactions—often varies from country to country.

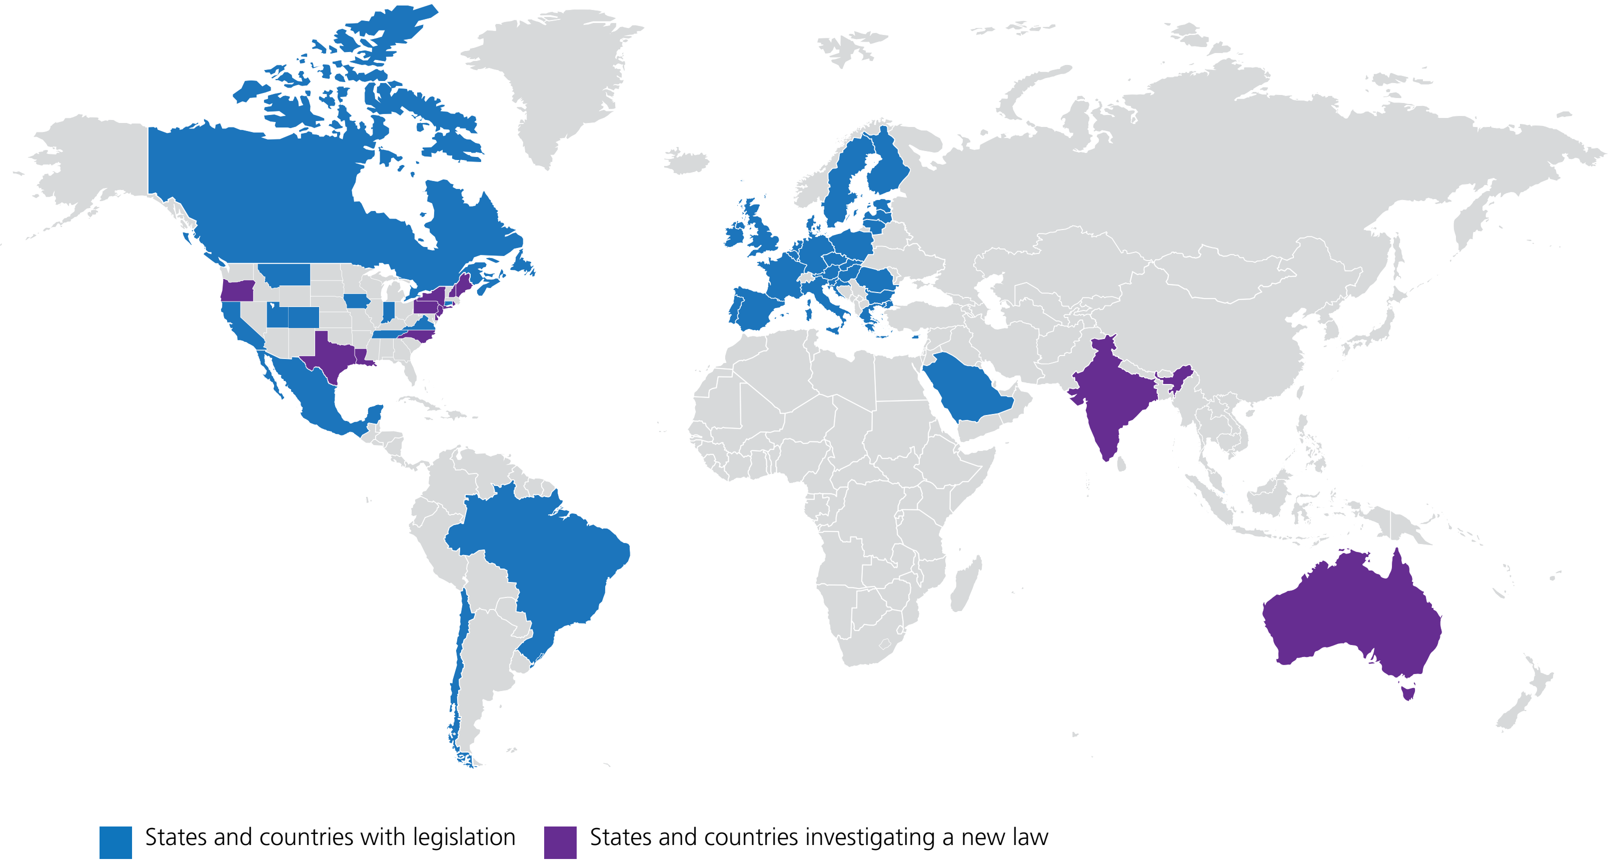
🟦 States and countries with legislation
🟪 States and countries investigating a new law

Figure 1: Data sovereignty legislations around the world

Around the world, many data sovereignty laws and regulations are already in place. Figure 1 illustrates a few examples.

Each set of regulations directly impacts the data centers operating within the state, region or country. For example, the EU is currently considering a plan that will force providers to store all their data within the bloc and require a cloud cybersecurity certification. Further, ENISA, the EU's cybersecurity regulator, is drawing up stricter requirements to ensure no foreign government can access EU data. Therefore, non-EU organizations may have to create "sovereign cloud" operations. Such cloud operations would need to be entirely located within the EU and compliant with EU rules that supersede all other national regulations. This could prove to be problematic for United States and Chinese tech companies. As an alternative, companies and EU-headquartered cloud providers are considering partnerships that could provide a workaround—at least temporarily.

Data Sovereignty vs. Data Residency

At first glance, the terms "data sovereignty" and "data residency" may seem closely related or even interchangeable. They are not.

Data sovereignty refers to the laws and governmental policies that apply to data stored in the country where it originated and is currently located. In light of the increasing globalization of the world's data and the rapid adoption of cloud systems, it is easy to understand the difficulties of enforcing and operating within the various data sovereignty guidelines.

Data residency refers to the decision of businesses to store data outside of the jurisdiction where it was created. Once the data is moved (and made available for storage or processing), it is subject to that specific region's laws, customs and expectations.

In other words, data sovereignty refers to the laws and governmental policies applicable to data stored in the country where it originated and is geographically located, while data residency refers to where the data is physically and geographically stored.

Coping with a Fractured System

For large multinational businesses, navigating the ins and outs of individual state and country data sovereignty regulations can significantly impede global growth. Decisions about where to locate a new facility can no longer be made purely on the strength of the business case; companies must also account for the local regulations and how they affect where data can be stored and processed. The following are just a few examples.

Data type and geography

The first consideration is the type of data to be stored and processed. For example, is it protected by personal privacy regulations (i.e., does it relate to a person's profile, employment, finance, health and/or payments)? Once the data type is characterized and understood, it must be evaluated within the context of the local or national data sovereignty laws. For example, some laws specify which type of data can or cannot leave the country of origin and cross national borders. Other laws allow moving some data types outside the country of origin, but only if the destination country has signed an equivalent privacy protection agreement (or law).

Finding the right data center fit

There are four choices available for hosting the data storage and processing, depending on the data type and geography:

- **On-premises**: In this scenario, data is processed and stored in an organization's own data center at a known location, which offers the best option for complying with data sovereignty regulations. An on-premises data center can also be designed to match the agile cloud performance needed to support advanced applications like machine learning and AI. Today, more on-premises data centers are being built with a cloud-first approach.
- **Hybrid cloud:** The hybrid cloud blends cloud-based efficiencies from anything-as-a-service (XaaS) providers with localized on-premises resources. On the one hand, the hybrid cloud model offers the flexibility, scalability, and cost structure of a large cloud provider, which makes it perfect for handling non-regulated data. On the other hand, data that must comply with local sovereignty regulations can be stored on-premises, enabling businesses to better manage diverse requirements.
- **Private cloud:** A private cloud involves the use of a massive cloud-based infrastructure—none of which is owned by the end user. However, the cloud provider can dedicate portions of the underlying IT infrastructure to a single customer and ensure customer access is entirely isolated. However, as with a hybrid cloud, the private cloud involves some tradeoffs. Having the IT infrastructure totally isolated provides the best opportunity to track and audit how the data is being stored and processed. Yet there is no guarantee that the data in a private cloud will comply with national or regional data sovereignty laws.
- **Public cloud:** A public cloud consists of masses of common IT infrastructure—none of which is owned by the end user. A public cloud houses data in off-premises data centers anywhere in the world, so the location and ownership of the data become non-issues.

Power and Location Matter

Once the type of hosting has been decided, it's critical to understand if the power required is available to support a new installation or an extension to an existing installation. The base unit of any data center is the server and racks that house the servers. On-premises hosting typically involves from 4 to 20 installed servers, each consuming approximately 1 kilowatt, depending on the processing required. A 100-rack installation in an on-premises data center with a power usage efficiency (PUE) of 1:2 could see a power draw of 1 to 5 megawatts.

For a cloud-scale deployment, the power draw becomes far greater, as the maximum server density will be required to support all operational models being offered by the cloud service provider. In this instance, it isn't unusual to have 25 servers per rack, with each rack using 20 to 80 kW of power and thousands of racks per location. Whether the data center is cloud-scale or on-premises, the location of compute and storage resources are directly impacted by the availability of power as well as data latency performance.

Physical Infrastructure Considerations

No matter where the data resides, the infrastructure must be built upon a strong yet agile passive cabling foundation. The physical layer infrastructure must be flexible enough to support the migration to higher data rates while satisfying the evolving requirements of the active equipment.

Fortunately, data center network topologies have evolved significantly, making supporting stable, flexible, future-ready deployments and applications easier and more efficient. One of the significant changes involves the migration from a three-tier approach (core, access and aggregation layers) to a Clos switching architecture, commonly referred to as "leaf-and-spine" (see Figure 2). This newer topology is based on an any-to-any connectivity approach that is ideally suited for today's high fiber-density designs. It offers flatter architecture with fewer "hops" between equipment, which enables easy expansion—with the only real limitation to horizontal expansion being the number of ports on the spine switches. Since the network is flatter and faster, the physical cabling infrastructure should be ready to support day one transmission speeds and future data rates.

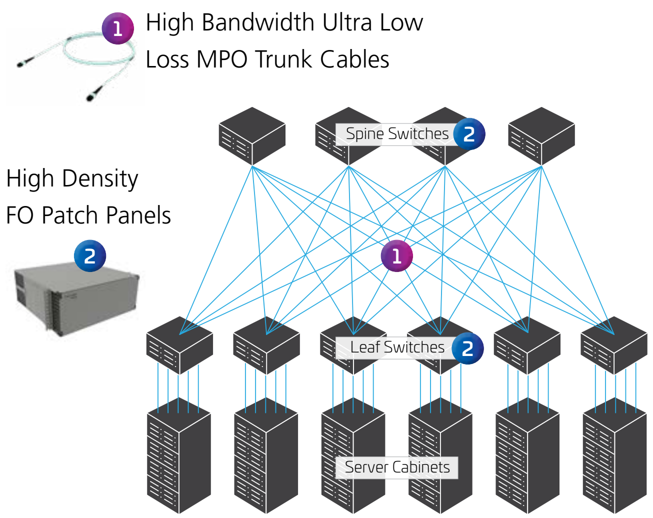

Figure 2: New flatter leaf-and-spine architecture using a three-tiered folded-Clos design for on-premises and hybrid data centers

In Figure 3, the stack on the left uses a traditional three-tier switching architecture, requiring the data to make six separate hops to traverse the switching layers and reach the destination server. The right-hand stack illustrates the inherent benefits of the flatter leaf-and-spine architecture. Reducing the number of switching layers decreases the number of hops—and associated latency—up to 33%.

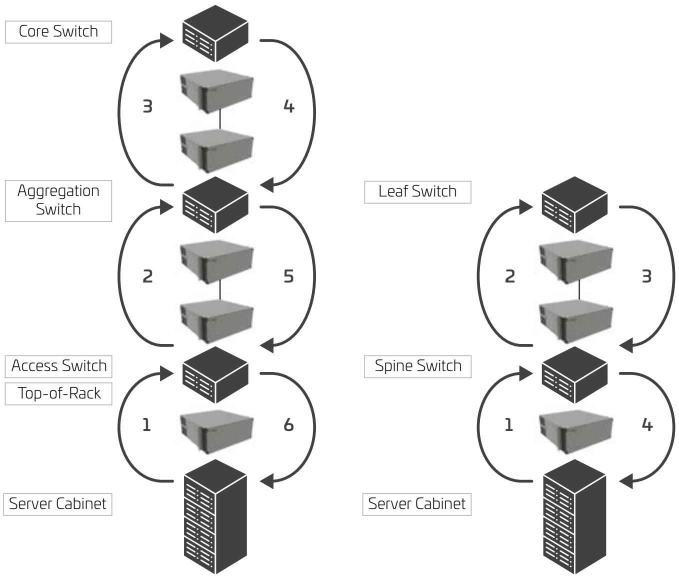

Figure 3: A traditional three-tier switching architecture versus a leaf-and-spine architecture

As the adoption of cloud-scale architectures continues to ramp up, the industry is coalescing around the leaf-and-spine architecture—with one small wrinkle: Satisfying the data handling requirements of much larger data centers is prompting many network managers to adopt a three-tier leaf-and-spine architecture, such as the one shown in Figure 4. These architectures will increasingly be supported by 16-fiber MPO connectivity as hyperscale and cloud-scale data centers migrate from 100G lane speeds to 400G, 800G, 1.6T and beyond. Figure 4 shows how 16-fiber MPO connectivity supports a three-tier Clos network.

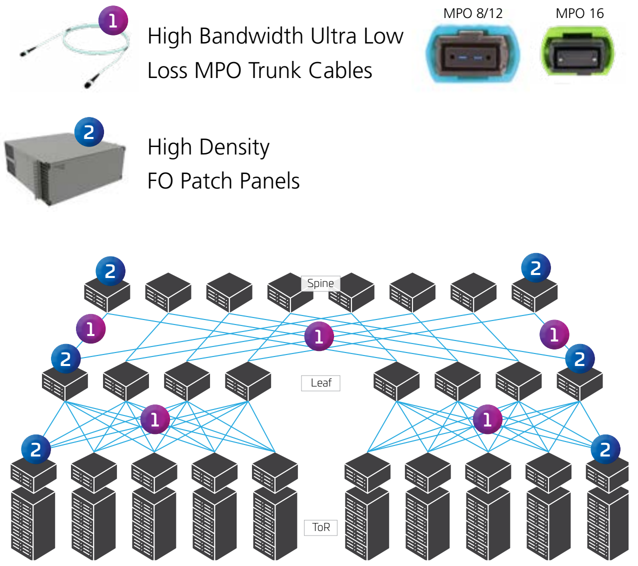

Figure 4: Three-tier Clos network with physical infrastructure components added, including MPO-16 connectors

Conclusion

While the definition of "data privacy" varies across regions, it is something that everyone agrees is critically important—particularly in an age of rapid network globalization. Looking into the future is always challenging, but if the past decade is any indicator, local variations in the generation and application of data sovereignty law will likely increase.

Building a data center infrastructure that supports data sovereignty is essential but not impossible. Success hinges on incorporating flexibility into the physical infrastructure to support future topologies and data rates for the next generation of servers and switches. With that in mind, the data center manager must be fluent in new connectivity technologies like MPO-16 and the fabric cabling needed to support growing bandwidth demands.

CommScope pushes the boundaries of communications technology with game-changing ideas and ground-breaking discoveries that spark profound human achievement. We collaborate with our customers and partners to design, create and build the world's most advanced networks. It is our passion and commitment to identify the next opportunity and realize a better tomorrow. Discover more at commscope.com.

commscope.com | Visit our website or contact your local CommScope representative for more information.