

Dynamic Path Selection

Digitization requires ever higher bandwidths, best network availability, and data security as well as fast time-to-service requirements. A software-defined WAN (SD-WAN) offers the best possible advantages for the networking of distributed company locations decisive advantages. It uses simple Internet-based wide-area networks (WAN) for rapid deployment of digital applications. When using multiple Internet lines at a single site, one of the greatest strengths of a modern SD-WAN comes into play: through “Dynamic Path Selection” (DPS), i. e. intelligent selection of the best quality line for applications, highest performance and maximum availability for digital applications is ensured. This techpaper illustrates the requirements, the functionality, and a practical example of how to implement Dynamic Path Selection in an SD-WAN.

The prerequisite

Active/active operation of multiple WAN lines

In order to meet the ever-increasing bandwidth requirements resulting from digitization, many companies rely on multiple parallel WAN lines in an SD-WAN. In this context, an SD-WAN is “line-agnostic”, which means that any type of line can be used flexibly and simultaneously—whether Ethernet, fiber, DSL / cable, external modem, or even cellular. One advantage of such active/active operations, i. e. the parallel use or Load Balancing of redundant Internet lines at a site, is the increase in the total bandwidth available. For example, by using two 300 Mbps DSL lines at the same time, a site has access to twice the bandwidth of 600 Mbps in total. Another advantage is the high availability of Internet access at the site: if one line fails, the affected sessions are routed via the active line(s). See “Figure 1: Active / Active link or Load Balancing”.




Figure 1:
Active / Active link
or Load Balancing

On this basis, an SD-WAN can distribute the load in the network to the individual lines according to the current requirements. This is where the “Dynamic Path Selection” technology comes into play.

How it works

Dynamic Path Selection

With Dynamic Path Selection, business-critical applications in an SD-WAN are always routed over the highest-quality line. All lines are continuously monitored using ICMP packets and metrics for load, packet loss, latency and jitter are calculated. The quality requirements of the individual applications are defined by predefined policies for lines (real-time video, voice over IP, low latency, best effort, and low priority), for example the allowed packet loss or the maximum latency of a line. The Dynamic Path Selection algorithm selects the line with the best quality for sessions of a specific application. If several lines meet the required policies, Load Balancing is performed in a round-robin process over these lines. Thus, users within large SD-WAN infrastructures with multiple WAN connections in active/active mode benefit from maximum performance and resilience. See “Figure 2: Functional principle of Dynamic Path Selection”.

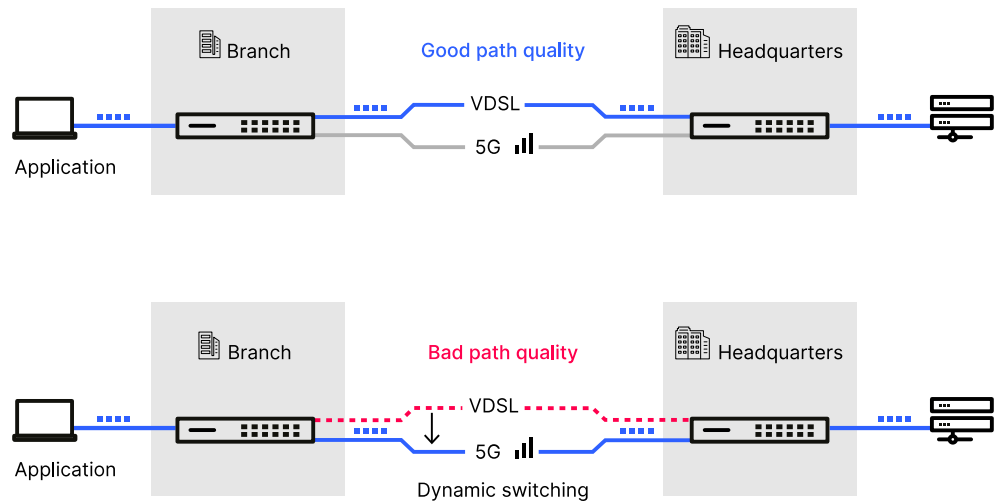


Figure 2:
Functional principle of
Dynamic Path Selection

Practical example

Scenario description

Dynamic Path Selection is designed to be used in classic SD-WAN scenarios. Such a scenario could look like this:

A branch office is connected to the headquarters via VPN. All Internet traffic is routed via the headquarters. Only trusted, per se encrypted applications, such as Office 365, are routed directly to the Internet via Local Internet Breakout, without detours via the headquarters. This reduces the Internet traffic on the central site and reduces the latency for the directly transmitted applications. For Load Balancing with Dynamic Path Selection to take place, the branch office has multiple WAN connections in place, ideally from different providers. These lines are active at the same time, i.e. in an “active / active” mode. The three lines are also defined as overlay networks, i.e. they are connected to the headquarters via VPN. See “Figure 3: SD-WAN scenario with Load Balancing, Local Internet Breakout, and Dynamic Path Selection”.




Figure 3:
SD-WAN scenario with
Load Balancing,
Local Internet Breakout,
and Dynamic Path Selection

Activation in the LANCOM Management Cloud

Make sure that the Internet connections are already configured on the branch gateways. You make this Internet connection known to the LMC in the **Sites > WAN connections** area. Dynamic Path Selection can now be activated by setting a check mark under **Project specifications > SDN > SD-WAN**. No further configuration steps are necessary, because policies, categories, and definitions for common applications and protocols are already predefined in the DPS configuration in the LANCOM Management Cloud (LMC). Thus, Dynamic Path Selection is applied immediately after activation and rollout of the

devices. Of course, additional applications can also be assigned quality criteria and individual settings can be made.

Monitoring in the LANCOM Management Cloud

The details view of the branch gateway in the Devices area contains the “Dynamic Path Selection” widget. This provides a high-level overview of the overall quality as well as the quality of the individual lines. A graphical interface allows the connection quality and the preferred lines for specific applications to be historically tracked. If only one WAN connection is available at the site, it is of course not possible to perform quality-based steering, but this can still be monitored in terms of quality using the Dynamic Path Selection widget.

The “All paths” graph shows the overall view of the quality across all lines. The higher the graph, the better the quality was at that point in time. In addition, the color scheme provides information about the line quality: If the graph is consistently green at the top, the overall quality across all lines meets the specifications for the selected category, even if individual lines were impaired in the meantime. On the time axis, you can also see via mouse-over which lines were the preferred ones at that point in time. On the right you can see the current state, i.e. which lines are the preferred ones and how many sessions are currently assigned to them.

In the left column you can switch between the different DPS categories, according to which the quality of the line is measured. The average quality of each line is displayed below. In the mouse-over of the lines per WAN connection, the recorded metrics for packet loss, latency, and jitter are displayed. In addition, the colored display gives an indication of the line quality at the respective point in time.

Therefore, if individual lines are impaired in practice and the impairments go beyond the expected behavior of this line, it is recommended to contact the respective provider. See “Figure 4: Monitoring widget “Dynamic Path Selection””.




Figure 4:
Monitoring widget
“Dynamic Path Selection”

Conclusion

The more bandwidth-intensive but also business-critical digital applications are used in a company, the more urgent is the need for automated monitoring of the quality of data transmission. Dynamic Path Selection is precisely this quality assurance in an SD-WAN. As shown by the example of the LANCOM Management Cloud, implementation is easy, and as a result all users in the network benefit from high connection quality at all times.