BROADCOM®

Product Brief

BCM65450

16-port G.fast Stackable Single-chip Solution

BCM65450

Key Features

- CO/CPE transceiver with integrated AFE for up to 20 lines
- 16 SerDes, with speeds ranging from 1G up to 25G, for flexible network interfacing (Ethernet with or without G.999.1), line card-level bonding, and vectoring
- Multimode and auto-mode support, selectable per channel: simultaneous support of G.fast, G.mgfast, and xDSL protocols on a per-line basis with embedded vectoring
- NTR, PPS, and ToD support (G.993.2 Amd7)
- 800 mW/channel in 16p G.fast 212a mode
- Advanced system-wide power management features, including integrated DC/DC AVS controller and RPF support
- 29 x 29 mm Flip Chip Ball Grid Array (FCBGA) with 1033 balls, 0.8 mm pitch, and embedded heat spreader
- 16 nm CMOS technology
- –40°C to +85°C ambient temperature operation

Overview

Broadcom's BCM65450 multimode DSL transceivers are the industry's most advanced and most flexible solution for DSLAM, MDU and DPU deployments. Both DSP and AFE functionalities are integrated into a single device.

The BCM65450 is a multichannel SoC supporting all VDSL, G.fast, and upcoming G.mgfast modes up to a 424 MHz bandwidth, and all VDSL profiles up to 35b.

The BCM65450 series is supported by multimode firmware fully compliant with ANSI T1.413 Issue I and II, ETSI TR328 and TS101-270, ITU-T G.992.1 (Annex A, B, C, and I), G.992.2 (G.lite), G.992.3 (Annex A, B, J, L, and M), G.992.5 (Annex A, B, J, and M), G.993.2 (G.vdsl2), G.993.5 (G.vector), G.9700/9701 (G.fast), G.996.1 (G.test), G997.1/2 (G.Ploam), G.998.1/2 (G.bond), G.998.4 (G.inp), and G.999.1 (G.int), as well as providing support for IEEE 802.3ah. Multiline CPE device mode is also available.

Applications

- Central Office Cabinet, MDU and DPU for Copper Access Networks
- 5G Backhaul

System Block Diagram

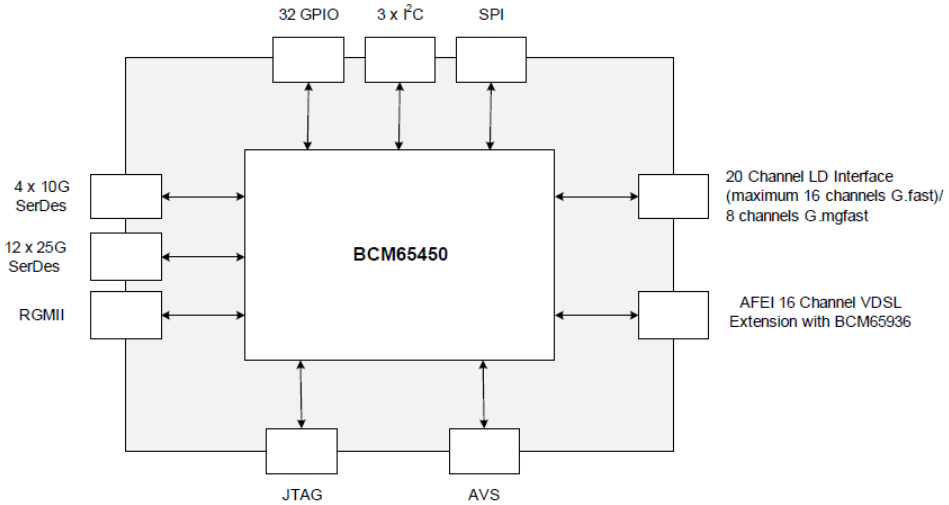

Broadband Solutions for Copper

As technically attractive as an all-fiber deployment may be, in almost all cases this is the most capital intensive and time-consuming option. With investors in telecom operators wary of substantial capital commitments, and with competitive and public pressure to improve service rates now, most operators are well-advised to wring the greatest possible advantage from their existing network. And with the latest innovations from Broadcom and industry standards bodies, they can now do so without any degradation in customer experience.

With the launch of the BCM65450 family, Broadcom builds on its 20-year history of delivering the best possible performance from the twisted pair copper plant. The 212 MHz iteration of G.fast offers usable customer speeds in excess of a gigabit per second, a fiber-like performance that provides a roadmap to higher bandwidth tiers that will be commercially competitive for many years into the future.

BCM65450

The BCM65450 offers a single chip 16-port solution for multi-gigabit services over copper. With 4x the density and 20% lower power than previous generations, the BCM65450 family enables the highest-density multi-dwelling unit (MDU) system designs. Each BCM65450 supports up to 16 G.fast 212 MHz interfaces, with embedded crosstalk cancellation (vectoring) for up to three devices interconnected seamlessly for a fully-vectored 48-port MDU. Coupled with Broadcom's previously announced BCM65550 vector processor, up to 192-port system designs are now possible. The BCM65450 also implements the latest ITU G.mgfast standard, which can deliver up to 8 Gbps by using 424 MHz of bandwidth. Additional information on newer features introduced with the BCM65450 is detailed below.

Symmetric Gigabit Service

In the drive to support Gigabit throughput for subscribers over copper infrastructure, operators today have had to use dynamic timeslot adjustment (DTA) so that the downstream or the upstream could run at 1 Gbps. Employing DTA meant that the other direction would simultaneously be skewed to be much less than 1 Gbps (on the order of 100 Mbps). DTA works best with only a few subscribers being coordinated. With more and more subscribers, the ability to skew in the downstream or upstream direction becomes more limited, often falling far short of a 1 Gbps connection. With the introduction of the BCM65450, operators can now offer 1 Gbps simultaneously in both directions. This full duplex (FDX) capability is possible with the introduction of echo cancellation techniques.

As shown in the diagram below, each side's transmission will be reflected, causing interference with the received signal. However, the transmitted signal is known, so the interference can be virtually eliminated by using echo cancellation (sometimes referred to as Self-NEXT). For twisted pair, far end crosstalk cancellation (FEXT) is also required, which is the traditional crosstalk cancellation for xDSL systems. One additional benefit of having a FDX system is that a packet can be acknowledged extremely quickly as data is always flowing in the opposite direction. The latency advantage of FDX over time division duplex (TDD is used for G.fast systems) results in a reduction from 1 ms to ~250 us. Note that FDX operation is possible with some lines still in TDD (single direction) mode without adversely affecting either the TDD lines or the FDX lines.

Echo Cancellation/Self-NEXT

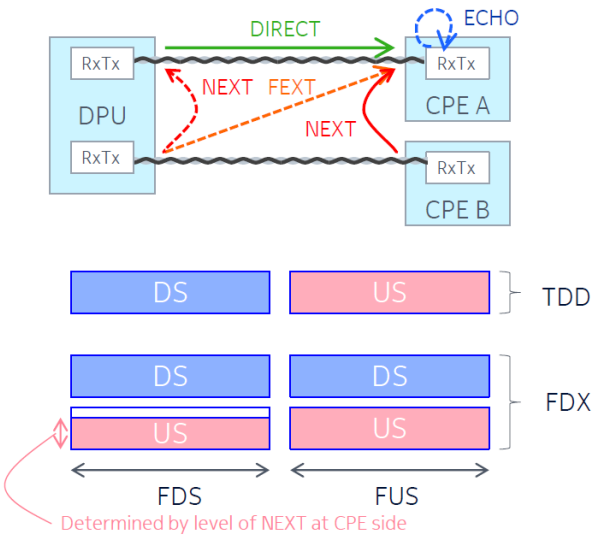

G.fastback

As operators continue to lay fiber further into their networks, there are still locations that are too far from the customer premise to enjoy gigabit speeds over the existing copper lines. The BCM65450 incorporates the new G.fastback ITU standard that enables bundling multiple copper lines into one big pipe, allowing a system to drop in between the fiber termination and the customer premise. This effectively halves the distance, enabling gigabit speeds for individual users as well as multi-gigabit uplinks for 5G cellular backhaul and other business opportunities. One key aspect of this functionality is it can be deployed over lines that are currently supporting POTS/ADSL/VDSL services, as well as over dedicated, unused lines.

G.fastback Deployment Options

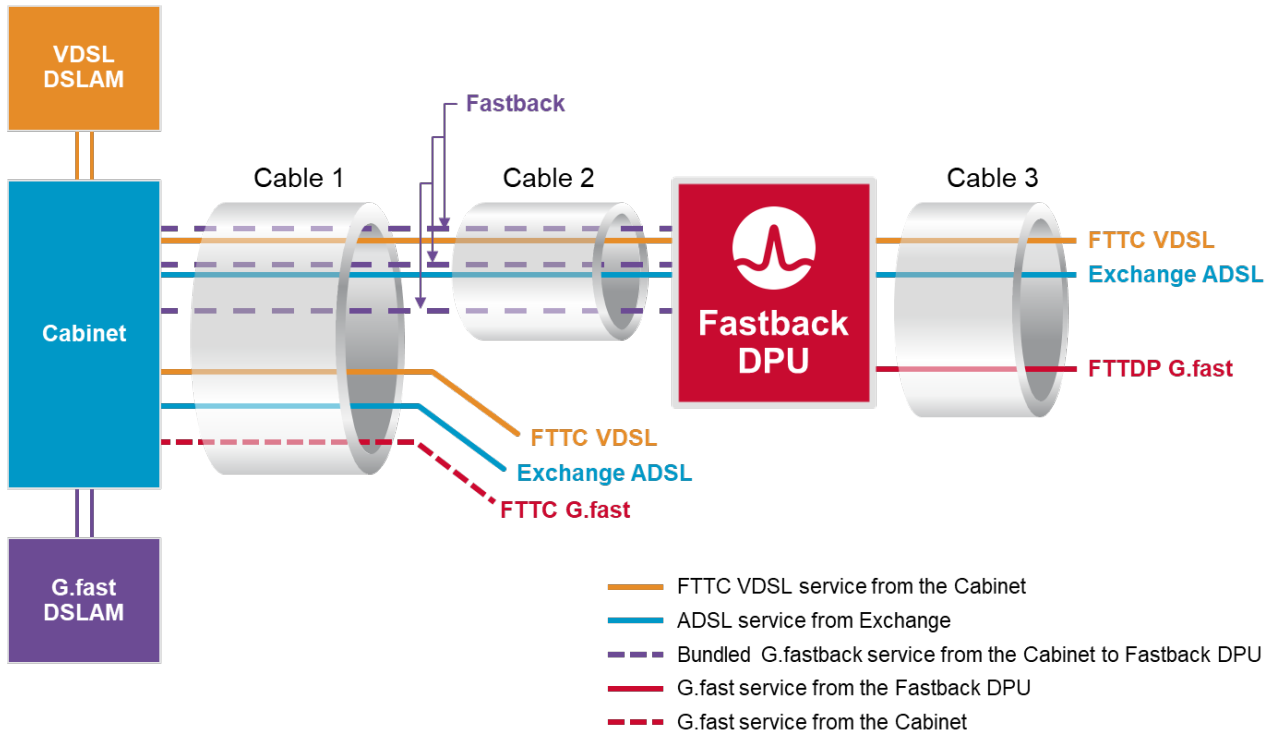

Cabinets are typically served by fiber links, which then connect to the copper network and provide local service to subscribers. This service is indicated by the three lines at the bottom of cable bundle 1. In many cases, this cable 1 bundle will also have services that come from the central office directly, and pass through the cabinet through a patch panel (top two solid lines in the cable 1 bundle), but still end up in the same cable bundle on the subscriber side of the cabinet with other lines being serviced by the fiber links. As can be seen in the drawing, these lines with service from the CO can be connected to the G.fastback DPU in the same manner, where the service can be passed through to the subscriber. G.fastback uses the bandwidth above the pass-through service (>2 MHz for ADSL service and >17 MHz for VDSL service) to supply additional services at half the distance, which is the red G.fast line to the right of the G.fastback DPU. This extra bandwidth is depicted by the purple dashed lines above the ADSL and VDSL lines to the G.fastback DPU. Note that any unused lines can also have this service, which is shown by the purple dashed line with no underlying service.

In order to use the extra bandwidth, the G.fastback DPU must be designed to handle additional crosstalk not normally seen at the cabinet. The reason is that when subscribers are transmitting on the uplink, the G.fastback DPU is receiving through the cable 3 bundle, and simultaneously transmitting on the cable 2 bundle. This causes crosstalk interference from the cable 2 bundle transmitting into the G.fastback receive of the cable 3 bundle. Similarly, when the cabinet is sending data on the downlink to subscribers, the G.fastback DPU is receiving through the cable 2 bundle, and simultaneously transmitting on the cable 3 bundle. In this case, crosstalk interference from the cable 3 bundle transmitting must be cancelled from the G.fastback cable 2 bundle receive. Note that in addition to this new crosstalk cancellation requirement, the two hop system architecture (cabinet to G.fastback DPU, and then G.fastback DPU to subscriber) adds a more significant delay than was specified in the original G.fast ITU standard. These concepts are captured in the diagram below.

G.fastback System Architecture

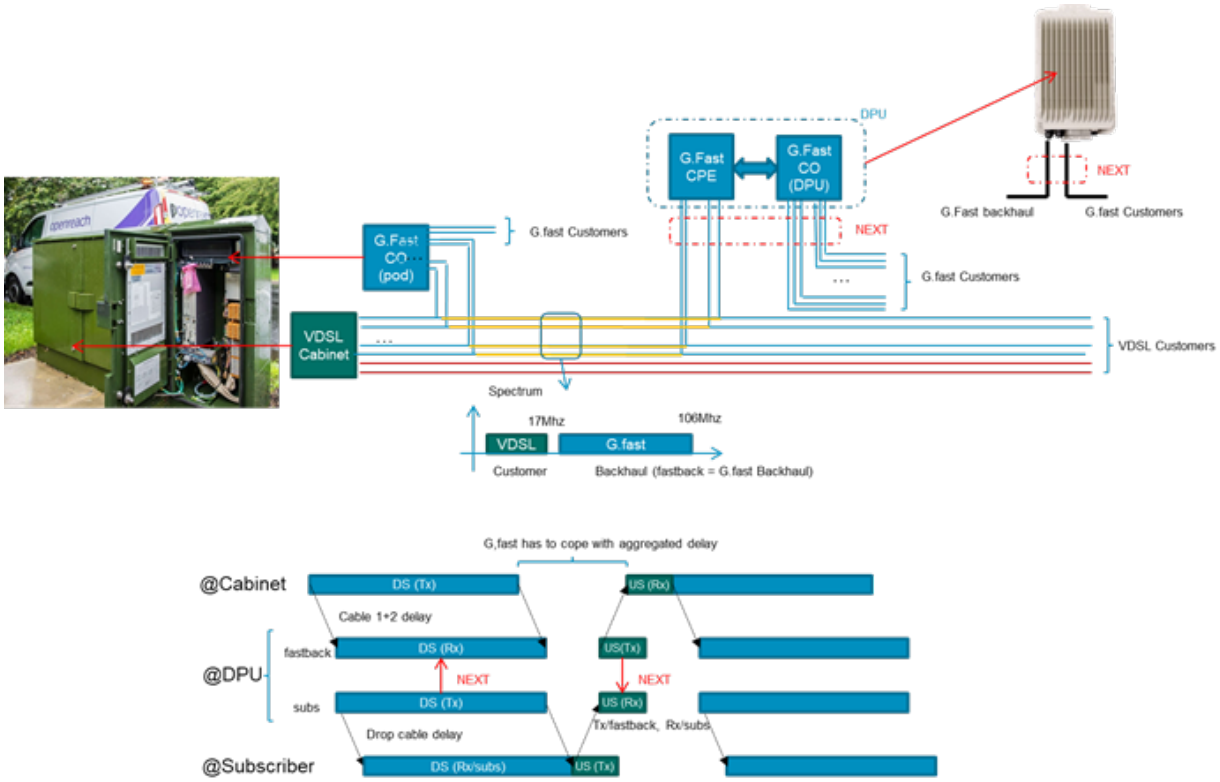

At the chip set level, Broadcom has implemented the upcoming G.fastback ITU standard using different individual chip sets from the Shasta family as depicted below. As noted in the diagram, eight G.fast 106 lines can be bundled on the uplink from the G.fastback DPU for an aggregated rate of up to 8 Gbps. On the downlink, up to 16 lines of G.fast 212 can be supported form the G.fastback DPU.

G.Fastback with BCM65400

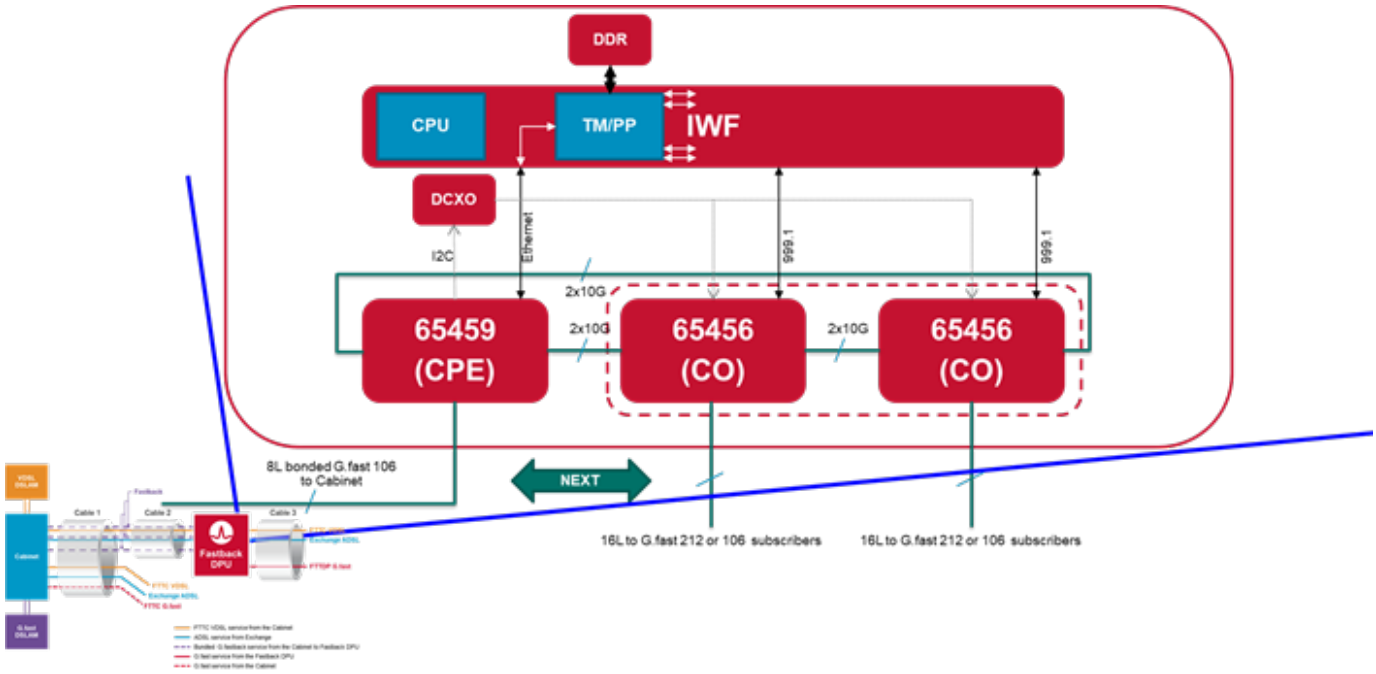

In some systems, all of the lines coming from the cabinet are providing G.fast service (there is no ADSL or VDSL services being provided). In these systems, the full bandwidth is being used, the 2 MHz or 17 MHz cut outs discussed above are not present. In these systems, there is no option for a bundled VDSL uplink, and the only option is for G.fastback. This implementation would essentially allow the G.fastback DPU to be a G.fast extender.

Bandwidth Boost

As discussed earlier, G.fast 212 systems deliver gigabit speeds asymmetrically by using DTA. The BCM65450 provides another tool for operators to use with these systems. Bandwidth Boost runs on top of G.fast 212 systems, allocating the 212 – 424 MHz band to one user, allowing that one user to get all of the additional throughput. This allocation between subscribers is done dynamically and can shift from one user to another as needed, and is independent in the downstream and upstream directions (one subscriber could get Bandwidth Boost in the downstream while another gets it in the upstream). The diagram below outlines Bandwidth Boost and how it can change over time.

Bandwidth Boost Over Time

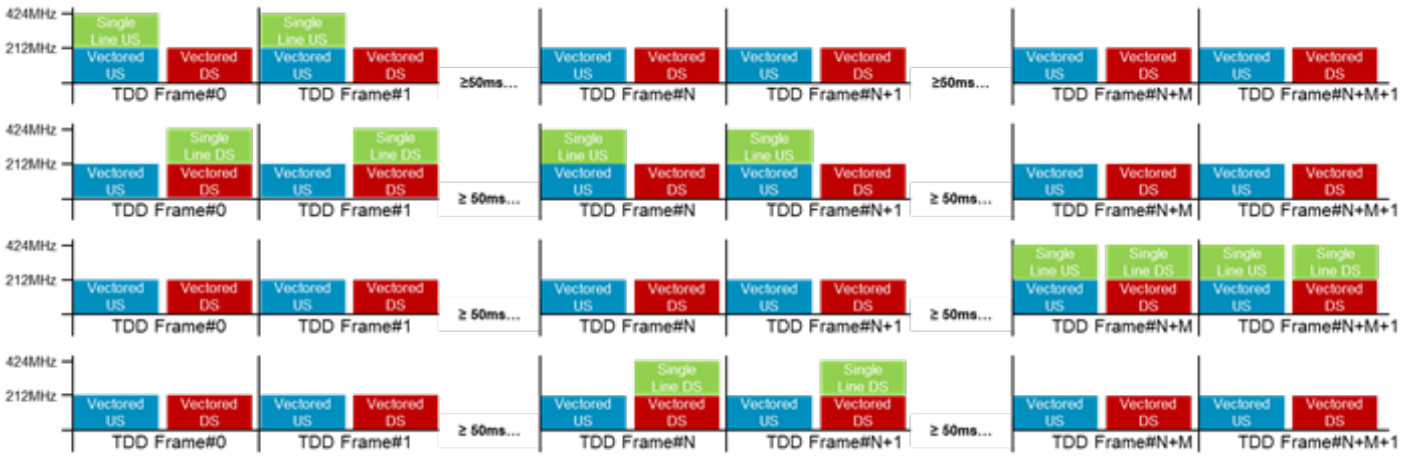

Point-to-Multipoint

G.mgfast runs up to 8 Gbps, a high enough rate that operators could entertain sharing this bandwidth among subscribers. The BCM65450 implements point-to-multipoint (P2MP) operation, which specifically address the ability to share this large throughput among up to four users. Two different scenarios for P2MP are envisioned: the first covers a traditional split among multiple homes, where more subscribers can be addressed from a single system. The second implementation is for use within certain types of homes that are constructed with materials that don't allow whole home WiFi coverage (for example, with walls made of stone or concrete). Bandwidth for each user is allocated dynamically, and can vary independently in the downstream and upstream directions (shown in the diagram below).

Point-to-multipoint (P2MP) Operation

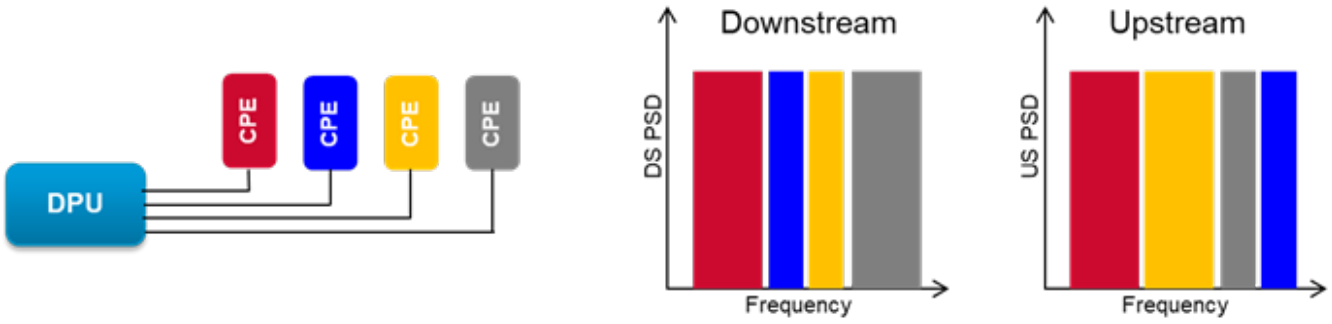

Ordering Information

Ordering Number	Description
BCM65450A0IFSBG	Product family code reference only, not orderable
BCM65452A0IFSBG	16p G.fast 106 (212-lite included) TDD/16p FDX coax/8p FDX TP
BCM65454A0IFSBG	4p G.fast 212 TDD/FDX coax/FDX TP
BCM65456A0IFSBG	16p G.fast 212 TDD/8p FDX coax/4p FDX TP
BCM65458A0IFSBG	8p G.fast 212 TDD/8p FDX coax/4p FDX TP
BCM65459A0IFSBG	8p G.fast 212 CPE (with bonding)
BCM65464A0IFSBG	4p G.fast 424 TDD/4p FDX coax/2p FDX TP
BCM65466A0IFSBG	16p G.fast 424 TDD PtMP
BCM65468A0IFSBG	8p G.fast 424 TDD/4p FDX coax/2p FDX TP
BCM65469A0IFSBG	4p G.fast 424/4p FDX coax/2p FDX TP CPE (with bonding)
BCM65475A0IFSBG	32p VDSL2 35b
BCM65476A0IFSBG	36p VDSL2 35b
BCM65479A0IFSBG	16p VDSL2 35b CPE (with bonding)

BROADCOM®
connecting everything®

For more product information: broadcom.com

65450-PB100 October 3, 2019