

D10056 DATACENTER SWITCH



Product Overview

The Inventec D10056 is a high performance and programmable ToR and aggregation switch designed for Data centers, large Enterprises as well as Service Provider network deployments.

The switch can be deployed in large scale layer 2 and layer 3 networks. Virtualized, overlay and traditional Data center networks are fully supported. Today's applications require networks to be Agile, Scalable, Flexible, Reliable, Programmable and Open.

The D10056 switch presents an open architecture with very high bandwidth and low latency design. It delivers 2.0Tbps throughput in a compact 1RU form factor. It offers 48x25GbE SFP28 ports and 8x100 GbE QSFP28 ports with PHY less design.

The SDN switch supports line rate L2/L3 forwarding, programmability, network virtualization, QoS and zero touch provisioning.

The D10056 offers customers a compelling choice between a high performance Quad core x86 or a very powerful Octal Core x86 CPUs based on functional and operational requirements.

Support for Open Network Ecosystems

The Inventec D10056 is an open switch that supports multiple Network Operating Systems (NOS). Today the switch ships as bare metal but can be integrated with Inventec SONiC or a third party NOS. The switch is SDN enabled and supports P4. Full ONIE support assures network operators of seamless integration into today's open network environments.

High Performance, Scalable and Flexible

The Inventec D10056 is a high performance switch allowing wire rate of 2B packets per second with a low cut-through mode latency, 22MB on-chip buffer memory and dynamic buffer management to prevent catastrophic TCP traffic scenarios. The switch offers scalability by supporting choice of high end X86 control processor with upto 16GB of fast DDR4 memory. With a PHY less design, the switch offers a low cut through latency.

Performance

- 2B Packets per Second
- 2.0 Terabits per Second Throughput
- Line Rate L2/L3 Forwarding
- 22 MB Packet Buffer
- 48x25 GbE SFP28 + 8x100GbE QSFP28

Control Plane

- CPU Options
 - 2.0 GHz x86 Octal-Core
 - > 8 GB to 16 GB DDR4 ^{NOTE1}
 - 2.2 GHz x86 Quad-Core
 - > 8 GB DDR4
 - 1.6 GHz x86 Quad-Core
 - > 8 GB DDR4 ^{NOTE1}
- 16 MB SPI Boot Flash with backup

Scalability

- 512K MAC Entries (HW Capable)
- 1.2M IPv4 Host Routes
- 1.3M IPv4 Routes
- 4 K VLANs

High Availability

- 1 + 1 Hot-Swappable & Redundant Power Supply
- 2 x SPI Flash Supports Boot Recovery
- 4 + 1 Hot-Swappable & Redundant Fans
- 802.3ad Link Aggregation/LACP
 - 256 Ports/Channel
 - 1024 Groups per System

Flexible Storage

- 8-128 GB SSD for Mass Storage
- 1x USB Port for External Storage

NOTE1 : These CPU variants are compatible but not fully tested. Will incur additional test time before shipping.



Inventec® at core

The switch is flexible and supports different cabling options as per customer needs. AOC (Active Optical Cabling) and pluggable transceiver optics of different lengths of fiber connections are supported. The port use is also flexible. Each QSFP28 100GbE port can be used as 4x25 GbE and all SFP28 25 GbE ports can be used in 10 GbE mode. The programmable parser allows for reconfigurability in the field and flexibility to cover a broad range of applications.

Agile, Programmable and supports Analytics

The switch supports a programmable parser with Match Action units. This allows for new protocol support. The switch supports a RESTful API interface. It allows for automatic provisioning and configuration with Puppet and Chef. Zero touch provisioning is also available. With Openstack orchestration tool integration, the switch enables automation, configuration and provisioning of L2 and L3 services in the data center. Switch also supports in-band Analytics and Telemetry support.

Rich Feature Set

The switch has a rich L2/L3 feature set to address the increase in datacenter network deployments and distributed computing applications. For cloud networking, it includes large L2/L3 switching & forwarding capacity and supports numerous multipathing and tunneling technologies and datacenter features like ECMP, VxLAN and NVGRE.

These overlays allow for network agility since the network operators do not have to modify the physical switch devices in case a user VM moves from one rack to another within the datacenter.

Available and Reliable Design

The switch is datacenter optimized with power and fan redundancy. It has a backup SPI boot flash that will activate for boot recovery if primary flash is corrupted. Also, with a PHYless switch, the overall MTBF increases with a less number of hardware components.

Applications

- Datacenter Spine Switch
- Service Provider Switch
- Large Enterprise

Programmability and Software Support

- Inventec INOS based on SONiC ^{NOTE2}
- ONIE
- Open Source Software Provided as RPM
- Chef and Puppet Client Integration
- Zero Touch Provisioning
- Bash Shell

Layer 2

- Dynamic ARP
- VLAN
- MTU Setting
- VLAN Trunk
- Static Port Breakout
- LAG
- MAC Ageing

Layer 3

- Address Resolution Protocol (ARP)
- Routing stack Graceful Restart
- L3 MLAG
- EVPN
- IPv6 link local
- IPv6
- ECMP
- LLDP
- QoS -ECN
- QoS-RDMA
- Border Gateway Protocol (BGP), Multiprotocol Extensions for BGP-4 (MP-BGP)
- BGP Graceful restart helper
- Fast Reload
- BGP/Neighbor-down fib-accelerate
- PFC WD
- Tunnel Decap

NOTE2 : The supported SONiC software features can be seen at this link below

<https://github.com/Azure/SONiC/wiki/Sonic-Roadmap-Planning>




Figure 1. Inventec Switches in a Leaf-Spine DC Architecture

In Summary

Overall performance, feature-richness, high availability, programmability, port-density, and line-rate switching capability makes the D10056 an excellent choice for next generation large and medium sized datacenters. This also makes the D10056 well suited for use as a general purpose switch in large Enterprises and Service Provider networks.

QoS

- Priority Flow Control
- COS

Network Management and Monitoring

- CLI
- Telnet/SSH
- Sflow
- Mirroring
- Critical Resource Monitoring
- SONiC to SONiC upgrade
- SNMP
- Syslog
- Sysdump
- Incremental Config
- DHCP Relay Agent
- NTP
- Sensor Transceiver Monitoring
- DHCP Relay Agent
- Multiple images support
- USB Port
- 1G Management ports

Security

- AAA (Accounting and Authorization)
- ACL Permit/Deny
- IPv6 ACL
- Dynamic ACL upgrade
- RADIUS
- TACACS+

Datacenter

- VRF
- Priority-based Flow control (802.1Qbb)
- L2 in L3 Tunneling (VxLAN/L2GRE/NVGRE)
- P4 Programming language Support

Category	Description	Specification
Physical	Form Factor	1RU Fixed
	Dimensions (D x W x H)	503.94 x 439.93 x 43.18 mm (19.84 x 17.32 x 1.7 inches)
	Weight	9.5 kg (20.9lbs)
	Interfaces	48 x 25GbE SFP28 and 8x100GbE QSFP28
	Power Supplies	2 (1+1) Hot-Swappable & Redundant
	Power Connector	IEC320-C13
	Fans	5 (4+1) Hot-Swappable & Redundant
	System Memory	8GB-16 GB
	Flash Storage	8-128 GB
	External I/O	1 x USB
	MGMT Port	1 x GE RJ-45
	Console Port	1 x RJ45 (RS-232)
	Reset	1 x Reset Button (Front Panel Mounted)
	Status LEDs	System Health Status/Fan status/Power Supply status
Activity LEDs	Link Activity/Status	
Optics and Cables	See Section "Supported Optics and Cables"	
Performance and Scalability	Forwarding	2 Bpps
	Throughput	2.0 Tbps Bi-directional
	Latency	650 ns
	Layer 2	512K Mac addresses, 4K Vlans
	Layer 3	1.2M IPv4 host routes, 1.3M IPv4 LPM
	Redundancy	256 x 802.3ad groups; 128-way ECMP
	Buffer	22 MB
	Memory	8 GB
Power	Type	AC OR DC
	Input Voltage	100~240 VAC / -40~72 VDC
	Input Frequency	50/60 Hz
	Typical/Max Power Draw	504 W/ 567 W
Cooling	Front to Back Airflow	Yes
	Back to Front Airflow	Yes
Environmental	Operating Temperature	0~40 °C
	Storage Temperature	-40~70 °C
	Relative Humidity	10~85%
	Altitude	0~3000 m (0~10,000 ft)
Compliance	EMI	CISPR-22/FCC Part 15 IEC61000-3-2/3 IEC61000-4-2/3/4/5/6/11
	Safety	CB: IEC60950-1 (2nd) CCC: GB 4943.1-2011
	RoHS	RoHS-6

Supported Optics & Cables

Speed	P/N	Type	Description
25G	LOHQF006-SD-R	DAC	0.5m SFP28
	MCP2M00-AXXX	DAC	0.5~3m SFP28
	RTXM330-003	AOC	3m SFP28 AOC
	RTXM330-030	AOC	30m SFP28 AOC
	MFA2P10-A005	AOC	5m SFP28 AOC
	FCBG125SD1C05 or FCCG125SD1C05	AOC	5m SFP28 AOC
	FCBG125SD1C30 or FCCG125SD1C30	AOC	30m SFP28 AOC
100G	RTXM420-550	SR4	MPO type 70m OM3,100m OM4
	FTLC1152RGPL	CWDM4	LC type 2km SMF
	LQ210PR-Oxxx	PSM4	MPO type 500m
	MMA1B00-C100D	SR4	MPO type 100m
	NDAQGF-0001	Fanout	1m fanout 25G 30AWG
	LOHQF004-SD-R	Fanout	3m fanout 25G 26AWG

ABOUT INVENTEC

Inventec Enterprise Business Group (EBG) was established in 1998 and has been focusing on the design and manufacturing of server systems. Inventec EBG is the key server system supplier of the global branding clients.

Network Infrastructure Design Center
 Inventec North America Corporation
 5201 Great America Pkwy., Suite 525
 Santa Clara, CA 95054
 Tel: +1-408-642-3395
 Email: switchinfo@inventec.com
 Website: <http://productline.inventec.com/switch/>



* Standards and RFC Compliance

RFC 1112 Host extensions for IP multicasting
 RFC 1256 ICMP router discovery messages
 RFC 1321 Message digest algorithm
 RFC 1519 CIDR
 RFC 1765 OSPF database overflow
 RFC 1812 Requirements for IPv4 routers
 RFC 1981 Path MTU for IPv6
 RFC 1997 BGP Communities Attribute
 RFC 2131 DHCP relay
 RFC 2236 IGMP v2
 RFC 2328 OSPFv2
 RFC 2365 Administratively scoped boundaries
 RFC 2370 The OSPF Opaque LSA Option
 RFC 2385 Protection of BGP Sessions via the TCP MD5 Signature Option
 RFC 2460 IPv6 Protocol Specification
 RFC 2461 Neighbor Discovery
 RFC 2462 Stateless Autoconfiguration
 RFC 2464 IPv6 over Ethernet
 RFC 2474 Definition of the differentiated services field (DS Field) in the IPv4 and IPv6 headers
 RFC 2475 An architecture for differentiated services
 RFC 2545 BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing
 RFC 2597 Assured forwarding PHB group
 RFC 2710 MLDv1
 RFC 2711 IPv6 Router Alert
 RFC 2918 Route Refresh Capability for BGP-4
 RFC 3021 Using 31 -Bit Prefixes on IPv4 Point-to-Point Links
 RFC 3046 DHCP/BOOTP relay
 RFC 3056 Connection of IPv6 Domains via IPv4 Clouds
 RFC 3101 The OSPF "Not So Stubby Area" (NSSA) option
 RFC 3137 OSPF Stub Router Advertisement
 RFC 3246 An expedited forwarding PHB (Per-Hop Behavior)
 RFC 3260 New terminology and clarifications for DiffServ
 RFC 3315 Dynamic Host Configuration Protocol for IPv6 (DHCPv6)
 RFC 3376 IGMPv3
 RFC 3484 Default Address Selection for IPv6
 RFC 3493 Basic Socket Interface for IPv6
 RFC 3513 Addressing Architecture for IPv6
 RFC 3542 Advanced Sockets API for IPv6
 RFC 3587 IPv6 Global Unicast Address Format
 RFC 3623 Graceful OSPF Restart
 RFC 3633 IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6
 RFC 3736 Stateless DHCPv6
 RFC 3768 Virtual Router Redundancy Protocol (VRRP)
 RFC 4213 Basic Transition Mechanisms for IPv6

RFC 4271 A Border Gateway Protocol 4 (BGP-4)
 RFC 4291 Addressing Architecture for IPv6
 RFC 4443 ICMPv6
 RFC 4456 BGP Route Reflectors
 RFC 4486 Subcodes for BGP Cease Notification Message
 RFC 4541 IGMP snooping
 RFC 4760 Multiprotocol Extensions for BGP-4
 RFC 5171 Unidirectional Link Detection (UDLD) Protocol
 RFC 5340 OSPF for IPv6
 RFC 5492 Capabilities Advertisement with BGP-4
 RFC 6164 Using 127-Bit IPv6 Prefixes on Inter-Router Links
 RFC 6583 Operational Neighbor Discovery Problems
 RFC 6860 Hiding Transit-Only networks in OSPF
 RFC 826 Ethernet ARP
 RFC 894 Transmission of IP datagrams over Ethernet networks
 RFC 896 Congestion control in IP/TCP networks
 RFC3810 MLDv2
 RFC3973 PIM-DM
 RFC4601 PIM-SM
 ANSI/TIA-1057 LLDP-MED
 Draft-ietf-idmr-dvmrp-v3-10 DVMRP
 Draft-ietf-magma-igmp-proxy-06.txt IGMP/MLD- based multicast forwarding (IGMP/MLD proxying)
 Draft-ietf-magma-igmpv3-and-routing-05.txt IGMPv3 and multicast routing protocol interaction
 IEEE 802.1AB Link level discovery protocol
 IEEE 802.1D Spanning tree
 IEEE 802.1p Ethernet priority with user provisioning and mapping
 IEEE 802.1Q Virtual LANs w/ port-based VLANs
 IEEE 802.1s Multiple spanning tree
 IEEE 802.1w Rapid spanning tree
 IEEE 802.1x Port-based authentication
 IEEE 802.3ac VLAN tagging
 IEEE 802.3ad Link aggregation
 IEEE 802.3x Flow control

Inventec® at core