

M6100 Chassis Switch

frequently asked questions



XCM8903SK, XCM8903SF, XCM8903SX, XCM8944, XCM8948, XCM8944F, XCM8924X, XCM89P, XCM89UP, APS1000W, AFT603

1. What is the NETGEAR® ProSAFE® M6100 Chassis Switch?

The NETGEAR ProSAFE M6100 Chassis Switch consists of Gigabit access layer and 10 Gigabit distribution layer switches in the NETGEAR modular chassis switch product line.

The M6100 switch series offers a high-quality, high-density chassis alternative to stackable switches in campus LAN and midsize organizations.

With more than 1.4 Tbps switching and routing capacity, passive backplane, management and fabric nonstop forwarding redundancy, the M6100 series delivers world-class resiliency and scalability.

Ultimately, operating software and system management features take the complexity out of delivering L2/L3/L4 rich services for enterprise edge and SMB core deployments.

3. Is there a License Upgrade for advanced new features?

No, the entire feature set is available without requiring any license. For simplicity, all advanced new features included in M6100 chassis series are available without additional cost.

4. What are the dimensions of the M6100-3S chassis?

The M6100-3S chassis provides extra high-density in 4U height (7 inches – 17.59cm) and 17.39 inches depth (44.16cm).

5. How many ports are supported?

The M6100-3S chassis supports a maximum of 144 RJ45 10/100/1000 ports, or 120 SFP 100/1000 ports, or 72 RJ45 10GBASE-T ports, or 48 SFP+ 1000/10GBASE-X ports, or a combination.

2. What are the available models?

SKU	DESCRIPTION
XCM8903SK	M6100-44G3-POE+ Starter Kit bundle including M6100-3S base chassis, XCM8944 blade, XCM89P daughter card, AFT603 fan tray, APS1000W power supply
XCM8903SF	M6100-44GF3 Starter Kit bundle including M6100-3S base chassis, XCM8944F blade, AFT603 fan tray, APS1000W power supply
XCM8903SX	M6100-24X3 Starter Kit bundle including M6100-3S base chassis, XCM8924X blade, AFT603 fan tray, APS1000W power supply
XCM8948	M6100 I/O Blade 48 x 1G (RJ45)
XCM8944	M6100 I/O Blade 40 x 1G (RJ45), 2 x SFP+ (independent), 2 x 10GBASE-T (RJ45, independent)
XCM8944F	M6100 I/O Blade 40 x 1G (SFP), 2 x SFP+ (independent), 2 x 10GBASE-T (RJ45, independent)
XCM8924X	M6100 I/O Blade 24 x 10GBASE-T (RJ45), 16 x SFP+ (shared with 10GBASE-T)
XCM89P	M6100 Daughter Card adding PoE+ 802.3at functionality to XCM8948 and XCM8944 blades (1 daughter card per blade required)
XCM89UP	M6100 Daughter Card adding UPOE functionality to XCM8948 and XCM8944 blades (1 daughter card per blade required)
APS1000W	PSU 1,000W AC (up to four PSUs in M6100-3S chassis; up to four more PSUs in RPS4000v2 external power supply bay)
AFT603	M6100 Fan Tray for M6100-3S chassis with front-to-back cooling principle

6. What are the new software features in the M6100 series?

• PV(R)STP (CLI only)
• CPU Rate Limiting enhancement
• Single Rate Policing (RFC 2697)
• PIM Timer Accuracy
• PIM-SM Unhandled Events
• OSPF ECMP Next Hop Rework
• OSPF Flooding Reduction
• IPv4 Fragmented packets
• MLAG
• BGP Route Ref (CLI only)
• Removing Private AS Numbers from a BGP Path
• OSPFv2/v3 Transit Network Hiding
• IPv6 enhancements
• BGP enhancements (CLI only)
• RSPAN – mirroring enhancements
• Policy-based Routing
• sFlow improvements
• Malicious Code Detection
• Track LAG Member Port Flaps
• MGMD Snooping
• Authentication Tiering DOT1X--MAB--Captive Portal
• OSPFv3 Enhancements
• DCBX (802.1Qaz) (CLI only)
• Priority Flow Control (PFC) (CLI only)
• Enhanced Transmission Selection (ETS) (CLI only)
• FCoE FIP Snooping (CLI only)

7. Why can't I find an ordering SKU for the empty M6100-3S chassis?

The M6100-3S empty chassis is not orderable as a separate SKU. The M6100-3S empty chassis only comes with M6100-44G3-POE+ or M6100-44GF3 or M6100-24X3 starter kit bundles.

The ordering SKU numbers are:

XCM8903SK-10000S (M6100-44G3-POE+ Starter Kit)

- One M6100-3S base chassis
- One I/O blade 40 x 1G (RJ45), 2 x SFP+, 2 x 10GBASE-T (XCM8944)
- One PoE+ daughter card (XCM89P)
- One fan tray (AFT603)
- One power supply (APS1000W)

XCM8903SF-10000S (M6100-44GF3 Starter Kit)

- One M6100-3S base chassis
- One I/O blade 40 x 1G (SFP), 2 x SFP+, 2 x 10GBASE-T (XCM8944F)
- One fan tray (AFT603)
- One power supply (APS1000W)

XCM8903SX-10000S (M6100-24X3 Starter Kit)

- One M6100-3S base chassis
- One I/O blade 24 x 10GBASE-T (RJ45), 16 x shared SFP+ (XCM8924X)
- One fan tray (AFT603)
- One power supply (APS1000W)

All these components are shipping in their own separate packaging and wrapped together when ordered under Starter Kit bundle SKU.

8. How can I rack mount the M6100-3S chassis?

The M6100-3S chassis comes with a rack-mount kit and complimentary handles for 2-post racks and wiring cabinets. The M6100-3S chassis also comes with sliding rails for 4-post racks and wiring/server cabinets.

9. Can any blade be inserted in any slot of M6100 base chassis?

Yes. Any of four M6100 I/O blades (XCM8948, XCM8944, XCM8944F, and XCM8924X) can be inserted into any slots of M6100-3S base chassis. Slot-1 is designed as a Supervisor slot. Slot-2 is designed as a Back-up Supervisor blade. Any combination of M6100 I/O blades (XCM8948, XCM8944, XCM8944F, and XCM8924X) can handle these roles.

10. Are the blades supporting hot-swap or hot-plug?

Yes, any M6100 I/O blades (XCM8948, XCM8944, XCM8944F, and XCM8924X) support hot-swap and hot-plug.

11. Why can't I find any Supervisory module SKU for fabric or management?

Because the M6100 series is designed upon an innovative "distributed fabric" which removes the need for a dedicated supervisory module, while providing a passive backplane, redundant fabric and redundant management – all with non-stop forwarding world-class resiliency and availability. As a result, all four M6100 I/O blades (XCM8948, XCM8944, XCM8944F, and XCM8924X) are equipped with dedicated hardware and software high-end distributed fabric on board. When inserted in Slot 1, any I/O blade will handle the Supervisory Role. When inserted in Slot 2, any I/O blade will handle the Back-up Supervisory role. It brings simplicity for management and maintenance or upgrades, while providing better availability and resiliency than any competitive solution in this price range.

12. What type of redundancy modes for fabric and management is available?

Only one: non-stop switching and hitless failover. M6100 series “distributed fabric” provides both Fabric and Management redundancy with non-stop switching and hitless failover. There is no fastidious configuration: Slot-1 is the Supervisor slot, and Slot-2 is the Back-up Supervisor slot (secondary). By default, the Back-up Supervisor slot (Slot-2) provides continuous stand-by mode within the distributed fabric. In case of a failure for the Supervisory blade in Slot-1, the Back-up Supervisor in Slot-2 is instantly taking over as the new Supervisor. Non-stop forwarding feature offers hitless failover for Slot-2 and Slot-3 switching and routing in that case.

13. What happens when Supervisory blade in Slot-1 gets replaced after removal?

When Supervisory blade in Slot-1 is removed, hitless failover automatically happens and the Back-up blade in Slot-2 becomes the ‘new’ Supervisor. Now, when the blade in Slot-1 has been serviced and/or replaced, nothing happens automatically, by design. The admin has to issue a command in the CLI (`movemanagement <2> <1>`) or perform a GUI configuration (`System > Chassis > Basic > Chassis Configuration`) in order to request reverse fail back to Slot-1 as the new Supervisor slot. This provides more control to the admin, while ensuring non-stop switching/routing and hitless fail back mechanism for the entire fabric and management. Network traffic isn't impacted by the fail back.

14. What type of redundancy modes for Power Supplies is available?

N+1 redundancy is provided when one power supply is in standby mode. There are four PSU slots in M6100-3S base chassis, so minimum two power supplies are required for N+1 mode. Four additional PSU slots can be added with RPS4000v2 external power supply bay. In that case, N+1 mode is available across the eight unified PSU slots.

15. Is M6100 chassis switch backplane ‘passive’?

Yes, M6100 backplane meets Passive Backplane requirements.

16. What is M6100-3S chassis fabric speed?

M6100-3S distributed fabric speed is 480 Gbps inter-module. Each slot provides 8 x 10G unidirectional access to the backplane (80G half-duplex and 160G full-duplex per slot).

17. What is M6100-3S routing and switching capacity?

M6100-3S chassis total routing / switching capacity is 1.4 Tbps intra-module. Each blade (line-card) provides local line-rate capacity.

18. What is M6100-3S total throughput?

M6100-3S chassis total throughput is 357 Mpps (Million packets per second) inter-module and 1,071 Mpps intra-module.

19. What is M6100-3S chassis PoE budget?

Depending on the number of power supplies and the I/O blades combination, M6100-3S chassis can provide up to 6,000W PoE budget. For instance, all 144 ports (Max 1G capacity) can provide each full PoE+ 30W per port or 41.5W UPOE power per port. Another example, up to 99 ports can provide full UPOE 60W per port.

20. How can I equip a Gigabit blade for PoE?

XCM8948 blade (48 x 1G RJ45) and XCM8944 blade (40 x 1G RJ45, 2 x SFP+, 2 x 10GBASE-T) can be equipped with a PoE daughter card for turning on PoE, PoE+ and/or UPOE on every Gigabit RJ45 ports on the blade. It means than any blade XCM8948 or XCM8944 can be upgraded with PoE at any point of time, and possibly downgraded if the application has changed. This is much more flexible than built-in PoE line-cards for instance, and it provides great investment protection in case of UPOE upgrade: only the daughter card would be upgraded, not the blade – hence lowering the cost and removing the need for re-configuration. There are two different PoE daughter cards: XCM89P for PoE+ and XCM89UP for UPOE. One daughter card per I/O blade is required.

21. Which PoE standards are supported by XCM89P and XCM89UP daughter cards?

XCM89P daughter card supports PoE (802.3af) and PoE+ (802.3at). XCM89UP daughter card supports PoE (802.3af), PoE+ (802.3at) and UPOE (Universal Power over Ethernet).

22. What is UPOE and what is supported exactly?

UPOE (Universal Power over Ethernet) is not an IEEE standard: UPOE was brought to market by Cisco® but its physical and LLDP specifications have been used by third-parties for interoperability, for instance via UPOE LLDP TLV, an 802.3 organizationally specific TLV. To power a UPOE PD device, all four pairs are used to deliver up to 30W + 30W = 60W per RJ45 Gigabit port. More generally, UPOE can deliver any PSE wattage up to 60W per port, when PoE+ is limited to 30W PSE maximum. When a XCM8948 blade (48 x 1G RJ45) or XCM8944 blade (40 x 1G RJ45, 2 x SFP+, 2 x 10GBASE-T) is equipped with XCM89UP UPOE daughter card, it is backward compatible with standard PoE (15.4W) and PoE+ (30W). This means that any standard PoE (802.3af) or PoE+ (802.3at) PD device is powered without any specific configuration. The NETGEAR M6100 chassis supports UPOE next-generation PD devices using UPOE LLDP TLV as the preferred method. M6100 chassis supports a second method for UPOE interoperability called forced 4-pair high power static method. A ‘PoE’ command must be issued in the CLI or the Web GUI to enable the forced 4-pair configuration. M6100 doesn't support Cisco CDP power negotiation mechanisms and algorithms for inline power negotiation with UPOE PD devices.

*Cisco® is a registered trademark of Cisco Technology, Inc.

23. What is PV(R)STP new feature?

This is Per VLAN Spanning Tree and Per VLAN Rapid Spanning Tree. There is no IEEE standard yet for Per VLAN (R)STP, but NETGEAR implementation follows the same rules than other vendor's Per VLAN (R)STP for interoperability:

- Per VLAN STP is similar to IEEE 802.1d (spanning tree protocol), but implemented per VLAN. A single instance of STP runs on each VLAN and each STP instance on a VLAN has a single root switch. FastUplink feature immediately moves an alternate port with lowest cost to forwarding state when the root port goes down to reduce the recovery time. It generates dummy SSTP frames to reduce convergence time. FastBackbonefast feature selects new indirect port when an indirect port fails.

• Per VLAN RSTP is similar to IEEE 802.1w (rapid STP), but implemented per VLAN. Per VLAN RSTP has in built support for FastUplink and FastBackbone, hence are not required.

On a given switch, only one mode is configurable at any instant, RSTP, STP, MSTP, PVSTP or PVRSTP. Enabling PVSTP or PVRSTP disables other spanning tree modes on the switch. STP (01:80:c2:00:00:00) and SSTP (shared spanning-tree protocol – 01:00:0c:cc:cc:cd) frames are generated for each STP instance when these features are enabled.

24. How FastUplink feature works?

The FastUplink feature is used for quick selection of a port with lowest cost, when the root port fails. In other words, FastUplink feature is used to reduce the time taken to converge, when a link fails. This feature is similar to Cisco® Uplink Fast feature. When the primary link fails, FastUplink will create an alternate path immediately. Since the ports that apply FastUplink don't have to wait for normal convergence time, this will speed up the transition from the failed primary link to the backup link (the port in a blocking state).

FastUplink feature is only for PVSTP. It is built-in for PVRSTP (no explicit configuration).

*Cisco® is a registered trademark of Cisco Technology, Inc.

25. How does the FastBackbone feature work?

The FastBackbone feature allows for faster convergence time when an indirect link to root fails. When a root port or blocked port receives an inferior BPDU from the designated switch on that port, the switch infers that an indirect link to the root has failed. To speed up convergence, the max age timer is immediately expired and the port is put through the Listening and Learning states.

Normal Timers =
Max Age (20 sec by default) + Listening (15 sec) + Learning (15 sec)

With FastBackbone =
Max Age (0 Expired) + Listening (15 sec) + Learning (15 sec)

FastBackbone feature is only for PVSTP. It is built-in for PVRSTP (no explicit configuration).

26. Is NETGEAR PV(R)STP interoperable with Cisco® PVST+?

Yes, NETGEAR PVSTP/PVRSTP and Cisco® PVST+/RPVST+ are interoperable.

PVSTP is similar to the MSTP protocol as defined by IEEE 802.1s, the main difference being PVSTP runs one instance per VLAN. In other words, each configured VLAN runs an independent instance of PVSTP. The protocol is equivalent to Cisco® PVST+ and can interoperate with each other.

PVRSTP is similar to the RSTP protocol as defined by IEEE 802.1w, the main difference being PVRSTP runs one instance per VLAN. In other words, each configured VLAN runs an independent instance of PVRSTP. Each PVRSTP instance elects a root bridge independent of the other. Hence there would be as many Root Bridges in the region as there are VLANs configured for PVRSTP. PVRSTP is equivalent to Cisco® RPVST+ protocol and can interoperate with each other.

*Cisco® is a registered trademark of Cisco Technology, Inc.

27. What is the difference between PV(R)STP and MSTP (Multiple Spanning Tree)?

The difference between MSTP and PV(R)STP is primarily in the way the protocol maps spanning tree instances to VLANs: PV(R)STP creates a spanning tree instance for every VLAN, whereas MSTP maps one or more VLANs to each MST instance.

28. What if there is a mismatch between one switch running PVSTP and another switch running PVRSTP?

PVSTP is based on 802.1s and PVRSTP is based on 802.1w. When a switch running PVRSTP receives PVSTP BPDU, then the switch falls back to the lower version (PVSTP) after its migration-delay timer expires.

29. How are BPDUs transmitted with PV(R)STP?

The switch running PV(R)STP transmits IEEE standard spanning tree BPDUs along with SSTP BPDUs. SSTP BPDUs are PV(R)STP dedicated BPDUs.

On Access Ports, the IEEE standard BPDUs and the SSTP BPDUs are transmitted as untagged packets. These ports are usually member of only one Access VLAN, and intend to connect devices that are not VLAN aware (untagged traffic on the device side).

On Trunk Ports, the IEEE standard BPDUs and the SSTP BPDUs are transmitted as untagged packets on the Native VLAN (PVID). For all other VLANs that are tagged, only SSTP BPDUs are transmitted, as tagged packets with respective VLANs.

30. How PV(R)STP interact with other IEEE standards STP/RSTP/MSTP?

A switch running PV(R)STP can communicate with a switch running IEEE standard spanning tree protocol.

To communicate each other, a switch running PV(R)STP will use the access VLAN instance on untagged ports (access ports), or the native VLAN (PVID) instance on tagged ports to send IEEE standard spanning tree BPDUs that are not tagged. If this switch receives IEEE spanning tree BPDU packets, it will include and process them in the access VLAN or native VLAN (PVID) instance on the port.

A switch running IEEE standard spanning tree protocol will use its CIST to communicate. If the switch receives the SSTP format BPDUs, it doesn't treat them as standard BPDUs. Hence the received SSTP formatted BPDUs are flooded on all the ports of the corresponding VLAN. The SSTP BPDUs are multicasted over the region. Only the untagged IEEE standard BPDU packets received from the switch running PV(R)STP are processed as standard BPDUs.

31. What is Single Rate Policing new feature?

Single Rate Policing feature enables support for Single Rate Policer as defined by RFC 2697:

- Committed Information Rate
 - Average allowable rate for the class
- Committed Burst Size
 - Maximum amount of contiguous packets for the class
- Excessive Burst Size
 - Additional burst size for the class
 - Credits refill at a slower rate than committed burst size
- DiffServ feature applied to class maps

32. What is Policy Based Routing new feature?

The Policy Based Routing feature overrides routing decision taken by the router and makes the packet to follow different actions based on a policy. It provides freedom over packet routing/forwarding instead of leaving the control to standard routing protocols based on L3.

For instance, some organizations would like to dictate paths instead of following the paths shown by routing protocols. Network Managers/ Administrators can set up policies such as:

- My network will not carry traffic from the Engineering department
- Traffic originating within my network with the following characteristics will take path A, while other traffic will take path B
- When load sharing needs to be done for the incoming traffic across multiple paths based on packet entities in the incoming traffic

Configuring PBR involves configuring a route-map with match and set commands and then applying the corresponding route-map to the interface. IP routing must be enabled globally and on the interfaces for PBR to operate. Policy Based Routing is applied to inbound traffic on interfaces.

37. What are the SFP/SFP+ optical transceivers supported by M6100 series?

NETGEAR GENUINE MODULES AND DAC CABLES:

MODEL	TYPE	DESCRIPTION	SFP PORTS	SFP+ PORTS
AFM735	100BASE-FX	NETGEAR SFP Multimode LC Transceiver	Yes	No
AGM731F	1000BASE-SX	NETGEAR SFP Multimode LC Transceiver	Yes	Yes
AGM732F	1000BASE-LX	NETGEAR SFP Single-mode LC Transceiver	Yes	Yes
AXM761	10GBASE-SR	NETGEAR SFP+ Multimode LC Transceiver	No	Yes
AXM762	10GBASE-LR	NETGEAR SFP+ Single-mode LC Transceiver	No	Yes
AXM763	10GBASE-LRM	NETGEAR SFP+ Multimode LC Transceiver	No	Yes
AXM764	10GBASE-LR LITE	NETGEAR SFP+ Single-mode LC Transceiver	No	Yes
AXC761	Direct-Attach	NETGEAR 1M SFP+ to SFP+ DAC Cable	No	Yes
AXC763	Direct-Attach	NETGEAR 3M SFP+ to SFP+ DAC Cable	No	Yes

Third party Direct-Attach Cables

M6100 Chassis Switch supports third-party DAC passive and active cables, as long as they are SFP+ MSA compliant.

Third party modules

ONLY THE FOLLOWING THIRD-PARTY MODULES ARE COMPATIBLE WITH M6100 CHASSIS SWITCH:

MODEL	TYPE	DESCRIPTION	SFP PORTS	SFP+ PORTS
J9099B	100BASE-BX10-D	HP® X112 100M SFP LC BX-D Transceiver	Yes	No
J9100B	100BASE-BX10-U	HP® X112 100M SFP LC BX-U Transceiver	Yes	No
J4860C	1000BASE-LH	HP® X121 1G SFP LC LH Transceiver	Yes	Yes
J9142B	1000BASE-BX10-D	HP® X122 1G SFP LC BX-D Transceiver	Yes	Yes
J9143B	1000BASE-BX10-U	HP® X122 1G SFP LC BX-U Transceiver	Yes	Yes
J9153A	10GBASE-ER	HP® X132 10G SFP+ LC ER Transceiver	Yes	Yes

*HP® is registered trademark of Hewlett-Packard Company. Above list shows that NETGEAR successfully tested interoperability between these HP® transceivers and M6100 chassis switch on the date of Oct, 1st 2014.

33. Does the M6100 series support dual images?

Yes, the M6100 series supports the Dual Image feature. The Dual Image feature allows for two images in the permanent storage. User can denote any of these images as an active image that will be loaded in subsequent reboot. This feature provides for reduced down-time for the switch, when the image is being upgraded/downgraded.

34. What is the USB port for on I/O blades front panel?

The USB port allows user to download/upload switch firmware or configuration file using USB flash device. It is also used to recover the firmware image through the utility menu during boot up. It is more effective and easier than using XMODEM serial port protocol for file transfer.

35. What are M6100 console ports for serial connection?

M6100 I/O blades provide two serial ports:

- One mini-USB console port
- One straight-through wiring RJ45 serial port

Both ports are active simultaneously. Mini-USB console port allows user to directly access M6100 switch using one USB cable. USB driver must be installed first. The USB driver can be obtained either from the CD that comes with the blade, either from the following link: <http://www.downloads.netgear.com/docs/m6100/enu/230-11407-01/index.htm>

Drivers for the mini-USB console port are provided for Windows Server 2008; Windows Server 2003; Windows 8; Windows 8 x64; Windows 7; Windows 7 x64; Windows Vista; Windows Vista x64; and Windows XP.

36. Is out-of-band management for Telnet, SSH and GUI network access supported by M6100?

Yes, M6100 Chassis Switch provides the admin with two differentiated methods (in-band and out-of-band) for Telnet, SSH and Web GUI network access. For security, the admin can decide to allow or restrict any of these two methods. Out-of-band management is possible through the dedicated OOB RJ45 10/100/1000 port on the front of any I/O blades (XCM8948, XCM8944, XCM8944F, and XCM8924X) when inserted in Supervisor Slot-1 (default). A second connection is recommended through OOB port on the blade inserted in Back-up Supervisor slot (Slot-2 by default) in case of failover. If OOB restriction is not a requirement, in-band management can be also available from any network port from any slot: Management ACLs are available to restrict which port(s) can reach M6100 CPU in that case.

38. How can I make sure I design and order M6100 components correctly, based on my needs?

You are more than welcome contacting NETGEAR Sales organization, for instance at (408) 907-8000 or by email: sales@netgear.com. NETGEAR Sales Engineers will help you design M6100 perfectly. You can also visit www.netgear.com/m6100 and download "M6100 basic configurator", or "M6100 expert configurator" files under Resources tab.

M6100 basic configurator: you can enter your port count needs and desired power requirements. A correct SKU list will automatically be provided.

M6100 expert configurator: if you are already familiar with M6100 series, you will select every blade one by one, and tailor power requirements based on your needs. A correct SKU list will automatically be provided.

39. What is the warranty of the M6100 Chassis Switch Series?

The M6100 series is covered under NETGEAR Lifetime Warranty and it includes:

- 90 days of Technical Support via phone and email
- Lifetime Technical Support through online chat
- Lifetime Next Business Day Hardware Replacement

You can find more information here: <http://www.netgear.com/business/documents/prosafe-lifetime-warranty/default.aspx> and here: <http://support.netgear.com/general/contact/default.aspx>

40. Where can I download software updates for M6100 Chassis Switch Series?

The M6100 series technical documentation and firmware updates can be found here: http://support.netgear.com/for_business/default.aspx

41. Where can I find more information on M6100 Chassis Switch Series?

Please visit www.netgear.com/managed and www.netgear.com/m6100