

**UNITED STATES DISTRICT COURT
FOR THE WESTERN DISTRICT OF TEXAS
WACO DIVISION**

INTELLECTUAL VENTURES I LLC
and
INTELLECTUAL VENTURES II LLC,

Plaintiff,

v.

LENOVO GROUP LIMITED,

Defendant.

Civil Action No. 6:23-cv-307

JURY TRIAL DEMANDED

COMPLAINT FOR PATENT INFRINGEMENT

Plaintiffs, Intellectual Ventures I LLC (“Intellectual Ventures I”) and Intellectual Ventures II LLC (“Intellectual Ventures II”) (together “IV”), for its complaint against Defendant Lenovo Group Limited (“LGL”) hereby allege:

THE PARTIES

1. Intellectual Ventures I is a Delaware limited liability company having its principal place of business located at 3150 139th Avenue SE, Bellevue, Washington 98005.
2. Intellectual Ventures II is a Delaware limited liability company having its principal place of business located at 3150 139th Avenue SE, Bellevue, Washington 98005.
3. Upon information and belief, LGL is a foreign corporation organized and existing under the laws of China, with its principal place of business located a No. 6 Chuang Ye Road, Haidian District, Shangdi Information Industry Base, Beijing, 10085, China and may be served pursuant to the provisions of the Hague Convention. Upon information and belief, LGL also has a principal place of business located at Lincoln House, 23rd Floor, Taikoo Place, 979 King’s Road, Quarry Bay, Hong Kong, S.A.R.

4. Upon information and belief, LGL is the parent company of a multinational conglomerate that operates under the name “Lenovo” and refers to itself and its subsidiaries as the “Group.” LGL purports to be a US\$60 billion Fortune Global 500 company serving customers in 180 markets around the world, including within the United States where it is a leading manufacturer and seller of laptop computers, desktop computers, smartphones, and tablets. Upon information and belief, each member of the “Group” is directly or indirectly a wholly owned subsidiary of LGL, which is the parent corporation or otherwise controls each member. These entities are collectively referred to herein as “Lenovo”.

5. Upon information and belief, LGL and each member of the “Group” are part of the same corporate structure and distribution chain and have acted in concert with respect to the facts alleged herein such that any act of LGL is attributable to every other member and vice versa.

NATURE OF THE ACTION, JURISDICTION, AND VENUE

6. IV brings this action for patent infringement pursuant to 35 U.S.C. § 271, *et seq.* This Court has subject matter jurisdiction under 28 U.S.C. §§ 1331 and 1338(a).

7. This Court has personal jurisdiction over LGL. LGL conducts business and has committed acts of direct and indirect infringement in this District, the State of Texas, and elsewhere in the United States. Moreover, LGL actively directs its activities to customers located in the State of Texas and this District.

8. For example, on information and belief Lenovo Group members, acting at the direction of LGL, sell within this District, the State of Texas and elsewhere in the United States products accused of infringement in this case.

9. Venue is proper in this District under 28 U.S.C. § 1391(c) because LGL is a foreign corporation. In addition, LGL has committed acts of patent infringement in this District and IV has suffered harm in this District.

FACTUAL BACKGROUND

10. Intellectual Ventures Management, LLC (“Intellectual Ventures”) was founded in 2000. Intellectual Ventures fosters inventions and facilitates the filing of patent applications for those inventions; collaborates with others to develop and patent inventions; and acquires and licenses patents from individual inventors, universities, corporations, and other institutions. A significant aspect of Intellectual Ventures’ business is managing the plaintiffs in this case, Intellectual Ventures I and Intellectual Ventures II.

11. One founder of Intellectual Ventures is Nathan Myhrvold, who worked at Microsoft from 1986 until 2000 in a variety of executive positions, culminating in his appointment as the company's first Chief Technology Officer (“CTO”) in 1996. While at Microsoft, Dr. Myhrvold founded Microsoft Research in 1991 and was one of the world’s foremost software experts. Between 1986 and 2000, Microsoft became the world’s largest technology company.

12. Under Dr. Myhrvold’s leadership, Intellectual Ventures acquired thousands of patents covering many important inventions of the Internet era, including many pertaining to the networked computers that comprise the Internet. Many of these inventions coincided with Dr. Myhrvold’s successful tenure at Microsoft.

Remote Network Device Management

13. One area of particular and continuing importance in the Internet era is the remote management of networked devices. Device security management specifically, which is the management of devices with the goal of protecting them from harm and unauthorized use, is

becoming more important with every passing year. Secure management of remotely located devices is essential for reliable, dependable, and highly available systems that are resilient to attack, responsive to customers' needs and affordable to operate.

14. Historically, the combination of remote management and secure management were not coextensive. As a result, networked devices were traditionally managed by physically isolating them, either individually or in small groups, from other parts of the network. An administrator for example, would typically co-locate several devices and limit physical access of those devices to select authorized employees. Any management of the devices would have to be performed while one of those employees was physically present with the devices. Such solutions became cost-prohibitive, in terms of both time and personnel, as networks grew and expanded over geographically dispersed areas.

15. When it became no longer feasible to have an administrator present at the location of every device in the network, many network administrators began allowing authorized employees to perform remote maintenance on networked devices. Enabling a device for remote management avoids the cost and delay of dispatching a person to the remote site, however, could potentially allow a determined intruder to utilize the remote access means for an attack if the remote management solution is not highly secure.

16. Remote management of network devices was performed over either "in-band" or "out-of-band" network connections. "In-band" management occurred over the same network that user data traversed, meaning that management data and user data flowed over the same network. "Out-of-band" management occurred using a means other than the network utilized for user data. Both "in-band" and "out-of-band" management did not have the appropriate level of security to prevent against potential attacks.

17. A disadvantage of out-of-band management arose because it bypassed several important network security systems that were employed by user data networks. These systems included virtual private networks (VPNs), firewalls, access control lists (ACLs) and authentication servers. As a result, out-of-band management made the network and its connected devices more vulnerable, in some ways, against malicious attacks.

18. In-band management also has its challenges. One is the comingling of user data and management data. Comingling of user and management data provides an opportunity for rogues to compromise management data from within the network itself, particularly if the administrator failed to implement a robust authentication scheme for other authorized administrators or employees. VPNs existed, which protected management data while it flowed over the in-band network, but even with VPNs, there was comingling of user data and management data in the device itself. In another example, existing authentication schemes, such as placing sole reliance on HTTPS authentication, were not always as robust as they needed to be. Problems such as comingling and authentication could be addressed by adding additional devices that implemented these features, which would be placed within or near the managed device. But such other devices would add cost and occupy extra space.

19. To overcome these obstacles, Engedi Technologies (“Engedi”), an early developer of network security solutions focused on secure remote management technology, and the original assignee of the patents-in-suit, developed the Secure Remote Management System (SRM). SRM provided an authenticated and encrypted secure tunnel between an SRM appliance co-located with a managed device, and a centralized network management center. These secure

network tunnels provided multi-pathed communication capability for the remote management of network devices.¹

20. SRM provided in-band and out-of-band secured network connections from the SRM appliance to the network management center, thus making available multiple and diverse robust paths for reporting status information to monitoring stations or allowing for remote configuration of the device. Compared to prior designs, this diverse and robust multi-path capability was a significant advantage.

21. Defendant makes, uses, and sells servers and network devices that include embedded secure management processors marketed under the Integrated Management Module II (“IMM2”) and XClarity Controller (“XCC”) brands (collectively “accused processors”), as well as purpose-built software that supports operation of the accused processors. These accused processors and purpose-built software are embedded in Defendant’s ThinkSystem, System x, and Flex System x, and BladeCenter servers, among others, as well as other solutions based on the aforementioned servers.

Microprocessor Clocking

22. Another area of particular and continuing importance in the computer era is that, as microprocessors grow in frequency and dimension, microprocessor clocks have become increasingly limited by wire delays. For example, the Pentium III microprocessor broke the 1 GHz barrier in 2000, and speeds have continued to increase ever since. At the same time, due to issues of reliability and performance, wire dimensions have been scaled in successive process generations

¹ For instance as described by Engedi at: <https://web.archive.org/web/20050309054746/http://www.engedi.net/focus.htm> and https://web.archive.org/web/20050130064915/http://www.engedi.net/documents/SecureRemoteManagement_ver2p5.pdf

more conservatively than transistor dimensions. The result of these frequency and dimensional trends is that microprocessor clock speeds became increasingly limited by wire delays, so much so that some of the subsequent microprocessors, e.g., the Pentium IV, had pipeline stages solely dedicated to moving clock signals across the chip.

23. Furthermore, a growing challenge has been to distribute the clock across a progressively larger die to increasing numbers of latches while meeting a decreasing clock skew budget. Researchers concluded that in order to continue the pace of clock frequency increases, microprocessor designers would be forced to abandon singly clocked globally synchronous systems in favor of some form of asynchrony.

24. Although purely asynchronous systems have the potential for higher performance and lower power compared to their synchronous counterparts, major corporations in the early 2000s were initially reluctant to fully migrate to asynchronous design methodologies. Two major reasons for this reluctance were the immaturity of asynchronous design tools relative to those in the synchronous domain, and the cost and risk of moving away from the mature design infrastructures rooted in synchronous systems that have been successfully used to create many generations of microprocessor products.

25. To address the issues identified above, David H. Albonesi led a team of professors and graduate students at the University of Rochester in developing a multiple clock domain microarchitecture that uses a globally asynchronous, locally synchronous clocking style.

26. Defendant makes, uses, and sells devices that include embedded ARM Cortex-Axx processors, as well as purpose-built software that supports operation of those processors.

Cyclic Diversity

27. A further area of importance in today's computing environments is that as wireless communications systems are widely deployed to provide various types of communications, demand for increased data rates has skyrocketed. This has led wireless system providers to develop new techniques for increasing data rates within the limited available radio frequency (RF) spectrum. One of these advancements has been the use of orthogonal frequency division multiplexing (OFDM) transmission. In OFDM transmissions, a radio channel is divided into a large number of closely spaced subchannels, an outgoing bitstream representing data to be transmitted is divided into multiple sub-bitstreams, and each sub-bitstream is transmitted over a subchannel in parallel with other sub-bitstreams that are each transmitted over their respective subchannels. Each such sub-bitstream is comprised of a series of symbols, that is, a waveform of the communication channel that persists for a fixed period of time, and from which data can be extracted by taking samples (i.e., measuring segments) of that waveform. Each of the symbols is separated by a guard interval (a gap in time between successive symbols that provides a buffer making transmission channels more resilient against the effects of a multipath propagation). The main advantages of OFDM is its ability to cope with severe channel conditions (e.g., signal fading, echoes, and interference).

28. As technology continued to advance, demand for increased speed and reduced interference resulted in the implementation of further improvements such as using multiple antennas in a single device, sometimes referred to as multiple input, multiple output (MIMO), which enables simultaneous or substantially simultaneous transmission of multiple bitstreams/sub-bitstreams in the same RF spectrum. When combined with OFDM, MIMO increases speed and improves reliability, however, it also introduces challenges, particularly when a multi-antenna MIMO enabled transmitter is communicating with a single antenna single input, single output

(SISO), receiver device. For example, signals transmitted from the MIMO transmitter may follow direct paths and multipaths to the SIS receiver, which can result in constructive interference (when multiple signals interact with one another to increase their amplitudes) or destructive interference (when multiple signals interact with one another to decrease their amplitudes), thus increasing packet error rates and causing other unwanted behavior that degraded the network quality.

29. One way that prior art systems addressed these inefficiencies was by implementing linear diversity schemes in which the transmission of one signal from a MIMO system is delayed relative to another signal from the MIMO system. Linear diversity schemes tend to reduce constructive and destructive interference by temporally decorrelating the transmissions of two signals, but they resulted in other problems such as one of the signals occupying the other's guard interval.

30. To address the inefficiencies set out above, cyclic diversity schemes were implemented (e.g., the cyclic-delay diversity scheme). In the cyclic-delay diversity scheme each of two or more transmitters send the same data in a respective stream of symbols, but cyclically offset one spatial stream vis-a-vis the other by a defined number of samples resulting in a circular shift of all the samples in a particular symbol (or part thereof). By introducing a relatively small cyclic delay to a first transmitted MIMO signal relative to a second transmitted MIMO signals, those of skill in the art were able to substantially reduce the problems set out above. But, by introducing a small cyclic delay between the first and second MIMO signals, upon receipt sometimes the receiver would be unable to determine whether the cyclic delay was intentional or caused by environmental or other factors. This inability in turn led to the receiver incorrectly assuming an attempt by the transmitter to beamform, which occurs when antennas are intentionally electronically steered to adjust the phase and amplitude of a transmitted signal at each antenna,

such that the signals combine constructively in the desired direction and destructively in other directions. That is, small cyclic delays were causing unintentional beamforming.

31. To address these and other problems in the art Mark Webster and Michael Seals, at the time engineers for Conexant Systems, developed improved systems and methods of wireless communication, which include, but are not limited to, an improved signal transmitting system capable of manipulating OFDM data packets and data streams using a cyclic diversity scheme based on cyclic advancement rather than cyclic delay, thereby improving packet reception performance and reducing packet error rates, among other benefits.

32. Defendant makes, uses, and sells devices that include embedded wireless 802.11n, 802.11ac and 802.11ax compliant chipsets configured to use MIMO and OFDMA with a cyclic shift diversity feature compliant with the respective 802.11 standard, such as the Motorola edge + smartphones.

Intra-cycle Timing Relationships in Integrated Circuits

33. An additional area of continued importance is the design and fabrication of high-performance signaling mechanisms for digital integrated circuit devices. For example, with respect to high-performance memory integrated circuit devices (e.g., Double Data Rate (“DDR”)) memory, ensuring the reliability in the design and fabrication of high-performance memory modules had become problematic for many OEMs by 2003. The slower memory bus speeds prior to 2003 had allowed significant specification margins in the design and fabrication of a given memory module. As memory bus speeds began approaching those supported by the Double Data Rate 2 (“DDR2”) standard in 2003 however, the industry recognized that their designs required ever more exacting control of critical timing specifications, and design parameters had to be even more strictly maintained to keep the entire system in balance. A stable DDR memory module had

to provide reliability, speed, and proper timing to ensure the overall system (e.g., CPU, bridge components, peripheral busses, etc.) operated at peak performance.

34. Transmeta Corporation (“Transmeta”), a fabless semiconductor company, was founded in 1995 to help address problems such as the ones identified above. Among other things, Transmeta developed low power x86 compatible microprocessors based on a very long instruction word (“VLIW”) core and a software layer called Code Morphing Software. Transmeta’s initial public offering on November 7, 2000, was the last of the great high-tech IPOs up through 2000, with its opening-day performance not being surpassed until Google Inc.’s IPO in 2004.

35. Transmeta launched its first product, the Crusoe processor, in January 2000, and its second processor, the Efficeon, in October 2003.

36. In October 2006, Transmeta sued Intel Corporation for infringement of ten Transmeta patents covering computer architecture and power efficiency technologies by making, using and selling a variety of microprocessors, including Intel’s Pentium III, Pentium 4, Pentium M, Core and Core 2 product lines. One year later, in October 2007, Transmeta and Intel settled the case, with Intel agreeing to pay Transmeta \$150,000,000 up front and an additional \$20,000,000 per year for five years.

37. Transmeta was acquired by Novafor Inc. in January 2009 for \$255,600,000. Intellectual Ventures acquired much of the patent portfolio—including one of the asserted patents discussed below—the following month.

38. Defendant makes, uses, and sells devices that include mobile phones, laptop computers and tablet devices that include Qualcomm processors and LPDDR4, LPDDR4X or LPDDR5 memory.

THE PATENTS-IN-SUIT

39. On January 29, 2008, the PTO issued United States Patent No. 7,325,140 (“the ’140 patent”), titled SECURE MANAGEMENT ACCESS CONTROL FOR COMPUTERS, EMBEDDED AND CARD EMBODIMENT.

40. The ’140 patent is valid and enforceable.

41. Intellectual Ventures II LLC is the owner and assignee of all rights, title, and interest in the ’140 patent, including the rights to grant licenses, to exclude others, and to recover past damages for infringement of that patent.

42. The ’140 patent is directed to a remote device management communication system including a secure management access controller embedded within and in direct communication with a managed networked device. The system includes at least one secure management access controller connected to one or more data buses, an out-of-band access connection for connecting network services or remote users with the secure management controller, at least one virtual management interface for connecting the network services or remote users with the secure management access controller, and where the virtual management interface connection provides logical separation of management data from user data and utilizes user interfaces of the managed network element for the connection to the one or more network services or remote users.

43. On June 25, 2013, the PTO issued United States Patent No. 8,474,016 (“the ’016 patent”), titled SECURE MANAGEMENT ACCESS CONTROL FOR COMPUTERS, EMBEDDED AND CARD EMBODIMENT.

44. The ’016 patent is valid and enforceable.

45. Intellectual Ventures II LLC is the owner and assignee of all rights, title, and interest in the ’016 patent, including the rights to grant licenses, to exclude others, and to recover past damages for infringement of that patent.

46. The '016 patent is directed to a computer network management apparatus and method for remotely and securely managing a networked device. The apparatus includes a processor configured to control one or more functions of a network device having a network interface, wherein the network device is configured to receive data requests and an encrypted form of management requests via the network interface, wherein the management requests are from a remote administrator; a first bus; a bus controller coupled to the processor via the first bus, wherein the bus controller is also coupled to a second bus of the network device that is distinct from the first bus; wherein the bus controller is configured to receive the encrypted form of the management requests from the second bus, and to convey the encrypted form of the management requests to the processor via the first bus; wherein the processor is configured to decrypt the encrypted form of the management requests; wherein the network device includes a processor configured to facilitate operation of the network device; and wherein the processor of the apparatus is distinct from the processor included in the network device.

47. The inventions claimed in the '140 and '016 patents were conceived by Jeffrey Alan Carley during his time as CTO and Co-Founder of Engedi. As noted above, Engedi created a secure remote management system to meet the need for a cost saving, highly secure method to access and manage remotely located devices in a distributed network. The system had a particular focus on preventing malicious attacks from insiders and resiliency in the event of path failures. Mr. Carley was an integral part of Engedi's technology development, architecting and overseeing the entire process, including managing funding, vendor and partner relationships and intellectual property growth. He has over 25 years of experience in the computer networking industry with major strengths in hybrid cloud networking, network architecture design and implementation, and network security and management at companies such as AIS, Pearson, TEKsystems, HPE, Modis,

MCI and IBM. Mr. Carley also holds the National Security Agency (NSA) InfoSec Assessment Management Methodology (IAM) certification and is a member of the IEEE, the Computer Society of the IEEE, the Information Systems Security Association and the Center for Internet Security. He is currently a Cloud Infrastructure Consultant at Applied Information Sciences and resides in Colorado Springs, Colorado.

48. On August 8, 2006, the United States Patent and Trademark Office issued United States Patent No. 7,089,443 (“the ’443 patent”), titled MULTIPLE CLOCK DOMAIN MICROPROCESSOR. The ’443 patent is valid and enforceable.

49. Intellectual Ventures I LLC is the exclusive licensee of the ’443 patent and has the right to sue and recover damages for any current or past infringement of the ’443 patent.

50. The ’443 patent is directed to a multiple clock domain (“MCD”) microarchitecture. In an MCD microprocessor, each functional block operates with a separately generated clock, while synchronizing circuits ensure reliable inter-domain communication.

51. The inventions claimed in the ’443 patent were conceived by David Albonesi, Greg Semeraro, Grigorios Magklis, Michael L. Scott, Rajeev Balasubramonian and Sandhya Dwarkadas at the University of Rochester. The first named inventor, Dr. Albonesi, is currently a full professor and the Associate Director of the School of Electrical and Computer Engineering at Cornell University, where he focuses on power-efficient computer architecture.

52. On November 24, 2009, the United States Patent and Trademark Office issued United States Patent No. 7,623,439 (“the ’439 patent”), titled CYCLIC DIVERSITY SYSTEMS AND METHODS. The ’439 patent is valid and enforceable.

53. Intellectual Ventures I LLC is the owner and assignee of all rights, title, and interest in the '439 patent, including the rights to grant licenses, to exclude others, and to recover past damages for infringement of that patent.

54. The '439 patent is directed to a system and method for transmitting OFDM signals from a multiple antenna transmitting device. The system is able to manipulate an OFDM signal using a cyclic advancement scheme whereby a portion of sampled symbol data from packets comprising the OFDM signal are shifted (advanced) into the guard interval of the packet relative to a first non-shifted version of the packet. The system and method then allow for the substantially simultaneous transmission of the respective packets from different antenna in the transmitting device, thereby allowing a receiver to more easily acquire and correlate the received data.

55. The inventions claimed in the '439 patent were conceived by Mark Webster and Michael Seals, both of whom were engineers at Conexant Systems, a well-known software developer and fabless semiconductor company specializing in developing technology for voice and audio processing. Mr. Webster is currently employed by L3Harris Technologies as a Senior Scientist, while Mr. Seals is a Principal Systems Engineer at Thales Group.

56. On January 12, 2010, the United States Patent and Trademark Office issued United States Patent No. 7,646,835 ("the '835 patent"), titled METHOD AND SYSTEM FOR AUTOMATICALLY CALIBRATING INTRA-CYCLE TIMING RELATIONSHIPS FOR SAMPLING SIGNALS FOR AN INTEGRATED CIRCUIT DEVICE. The '835 patent is valid and enforceable.

57. Intellectual Ventures II LLC is the owner and assignee of all rights, title, and interest in the '835 patent, including the rights to grant licenses, to exclude others, and to recover past damages for infringement of that patent.

58. The '835 patent is directed to the automatic calibration of intra-cycle timing relationships between command signals, data signals, and sampling signals for an integrated circuit device, including the generation of command signals to access an integrated circuit component; the accessing of data signals to convey data for the integrated circuit component; the accessing of sampling signals to control sampling of the data signals; and systematically altering a phase shift of the command signals, a phase shift of the data signals, and a phase shift of the sampling signals to determine a valid operation range of the integrated circuit device, wherein the valid operation range includes an optimal operation point for the integrated circuit device.

59. The inventions claimed in the '835 patent were conceived by Guillermo J. Rozas during his time at Transmeta Corporation.

COUNT I

(Defendant's Infringement of U.S. Patent No. 7,325,140)

60. The preceding paragraphs are reincorporated by reference as if fully set forth herein.

61. The '140 patent claims and teaches, *inter alia*, an improved way to provide secure remote management for devices by deeply embedding the necessary secure remote management hardware and software in the managed device itself. The inventions improved upon then-existing remote access/management security techniques by combining such hardware with a virtual management interface for logically separating user data from management data when using in-band management techniques. They added critical features to in-band management such as enabling separation of management and user data when administrators used in-band management all the way up to the network port itself. They accomplished this by creating a virtual interface at the physical port that accepts management and user data to keep the two data types segregated from end to end, including within the managed device, and not just on the network. Furthermore,

this was accomplished without requiring adding more devices in or around the managed device by embedding the secure remote management hardware and software into the managed device itself. This realized significant costs savings for customers that otherwise would have had to add more devices that took up more space substantially increasing cost. Further security and redundancy improvements were provided by the establishment of a separate purpose-built network connection interface for the secure remote management of the device over an out-of-band connection.

62. More specifically, the claims of the '140 patent recite a remote device management communication system for securely controlling access to management applications and communications to and from said management applications on network devices in a distributed computer network. The system includes one or more network services, one or more secure management access controllers, and one or more managed network devices. Further, the system includes at least one secure management access controller connected to one or more data buses of the managed network device for communication of device management data, as well as an out-of-band access connection means for connecting one or more network services or remote users with the secure management access controller for management of the network device. In addition, the system includes at least one virtual management interface connection means for connecting said one or more network services or remote users with the secure management access controller, where the virtual management connection means provides logical separation of management data from user data and utilizes user interfaces of said managed network element for connecting said one or more network services or remote users with the secure management access controller.

63. The system covered by the asserted claims, therefore, differs markedly from the prior systems in use at the time of this invention, which lacked the claimed combination of the secure management access controller connected to a managed network device, an out-of-band

connection means for connecting one or more network services or remote users with the secure management access controller, and a virtual management interface connection means for providing logical separation of management data and user data and for utilizing user interfaces of the managed device to also connect one or more network services or remote users with the secure management access controller.

64. Defendant has directly infringed, and continues to directly infringe at least claim 1 of the '140 patent by making, using, testing, selling, offering for sale, and importing into the United States products and services covered by one or more claims of the '140 patent. Defendant's products and services that infringe the '140 patent include all products and services that use an IMM2 or XCC embedded processor, which include the System x, NeXtScale, Flex System, BladeCenter, and ThinkSystem servers series, and any chassis/enclosures such products may be housed in, as well as any other of Defendant's products and services, either alone or in combination, that operate in substantially the same manner (together, the "Accused '140 Products").

65. Claim 1 of the '140 patent is reproduced below:

1. A remote device management communication system for securely controlling access to management applications and communications to and from said management applications on network devices in a distributed computer network that includes one or more network services, one or more secure management access controllers, and one or more managed network devices, the remote device management system comprising:

at least one secure management access controller connected to one or more data bus of said managed network device for the communication of device management data;

an out-of-band access connection means for connecting said one or more network services or remote users with said secure management access controller for management of said network device; and

at least one virtual management interface connection means for connecting said one or more network services or remote users with said secure management access controller;

wherein said virtual management interface connection means provides logical separation of management data from user data and utilizes user interfaces of said managed network element for connecting said one or more network services or remote users with said secure management access controller.

66. The Accused '140 Products each provide a remote device management communication system for securely controlling access to management applications and communications to and from said management applications on network devices in a distributed computer network that includes one or more network services, one or more secure management access controllers, and one or more managed network devices. As one example, the Accused '140 Products are network devices, modules and nodes capable of being configured in a distributed computer network, such as the Lenovo ThinkSystem SD650-N V2, ThinkSystem SN550 V2, and ThinkSystem SN550 (Xeon SP Gen 1) servers that support the XCC Standard, SXX Advanced and XCC Enterprise integrated service processors for controlling and securing remote management applications and services as well as communications regarding the same, as seen below:

Lenovo XClarity Controller

The Lenovo XClarity Controller (XCC) is the next generation management controller for Lenovo ThinkSystem servers. Language:

- [Lenovo XClarity Controller with Intel Xeon SP \(1st, 2nd Gen\)](#)
- [Lenovo XClarity Controller with Intel Xeon SP \(3rd Gen\) and AMD EPYC \(2nd, 3rd Gen\)](#)
- [Lenovo XClarity Controller REST API reference](#)

The following table lists the server models supported by LXCC Products:

Table 1. Server models supported by LXCC

Lenovo XClarity Controller	Server models	
Lenovo XClarity Controller with Intel Xeon SP (1st, 2nd Gen)	<ul style="list-style-type: none"> • SD530 • SD650 • SE350 • SN550 • SN850 • ST250 • ST258 • ST550 • ST558 	<ul style="list-style-type: none"> • SR570 • SR590 • SR630 • SR650 • SR670 • SR850 • SR850p • SR860 • SR950
Lenovo XClarity Controller with Intel Xeon SP (3rd Gen) and AMD EPYC (2nd, 3rd Gen)	<ul style="list-style-type: none"> • SD630 V2 • SD650 V2 • SD650-N V2 • SN550 V2 • ST650 V2 	<ul style="list-style-type: none"> • SR630 V2 • SR645 • SR650 V2 • SR665 • SR670 V2 • SR850 V2 • SR860 V2

Source:

https://sysmgt.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.lxca.doc%2Fremotecontrol_thinksystem_use.html

Server support

The following table shows what level of XCC is included with each ThinkSystem server.

Table 3. Server support

Server	XCC Standard	XCC Advanced	XCC Enterprise
ThinkSystem V2 servers			
ST650 V2 (7Z74/7Z75)	Supported	Supported	Supported
SR630 V2 (7Z70/7Z71)	Supported	Supported	Supported
SR650 V2 (7Z72/7Z73)	Supported	Supported	Supported
SR670 V2 (7Z22/7Z23)	Supported	Upgrade	Upgrade
SR850 V2 (7D31 / 7D32 / 7D33)	N/A	N/A	All models
SR860 V2 (7Z59 / 7Z60)	N/A	N/A	All models
SD630 V2 (7D1K)	Supported	Upgrade	Upgrade
SD650 V2 (7D1M)	Supported	Upgrade	Upgrade
SD650-N V2 (7D1N)	Supported	Upgrade	Upgrade
SN550 V2 (7Z69)	N/A	N/A	All models
ThinkSystem V1 servers			
SE350 (7Z46 / 7D1X)	Supported	Most models**	Some models**
ST50 (7Y48/7Y50)	Not supported	Not supported	Not supported
ST250 (7Y45/7Y46)	Most models*	Upgrade	Upgrade
SR150 (7Y54)	Most models*	Upgrade	Upgrade
SR250 (7Y51/7Y52)	Most models*	Upgrade	Upgrade
ST550 (7X09 / 7X10)	Most models*	Upgrade	Upgrade
SR530 (7X07 / 7X08)	Most models*	Upgrade	Upgrade
SR550 (7X03 / 7X04)	Most models*	Upgrade	Upgrade
SR570 (7Y02 / 7Y03)	Most models*	Upgrade	Upgrade
SR590 (7X98 / 7X99)	Most models*	Upgrade	Upgrade
SR630 (7X01 / 7X02)	Most models*	Upgrade	Upgrade
SR635 (7Y98 / 7Y99)	Not supported	Not supported	Not supported
SR645 (7D2Y/7D2X)	Most models*	Upgrade	Upgrade
SR650 (7X05 / 7X06)	Most models*	Upgrade	Upgrade
SR655 (7Y00 / 7Z01)	Not supported	Not supported	Not supported
SR665 (7D2W/7D2V)	Most models*	Upgrade	Upgrade
SR670 (7Y36 / 7Y37 / 7Y38)	Configure-to-order	Configure-to-order	Configure-to-order
SR850 (7X18 / 7X19)	N/A	Most models**	Upgrade
SR850P (7D2F / 2D2G)	N/A	N/A	Most models**
SR860 (7X69 / 7X70)	N/A	Most models**	Upgrade
SR950 (7X11 / 7X12 / 7X13)	N/A	N/A	All models
SD530 (7X21)	Most models*	Upgrade	Upgrade
SD650 (7X58)	Configure-to-order	Configure-to-order	Configure-to-order
SN550 (7X16)	N/A	N/A	All models
SN850 (7X15)	N/A	N/A	All models

Part numbers

Models of ThinkSystem servers come with either XClarity Controller Standard, Advanced or Enterprise, depending on the server type and the model. The servers will be delivered with the stated version already active. The following table shows the field upgrades available for models that come with XCC Standard or XCC Advanced.

Source: <https://lenovopress.com/lp0880-xcc-support-on-thinksystem-servers>

Lenovo ThinkSystem SD650-N V2 Server


Manageability and security

The following powerful systems management features simplify local and remote management of the SD650-N V2 server:

- The server includes an XClarity Controller (XCC) to monitor server availability. Optional upgrade to XCC Advanced to provide remote control (keyboard video mouse) functions. Optional upgrade to XCC Enterprise enables the additional support for the mounting of remote media files (ISO and IMG image files), boot capture, and power capping.
- Lenovo XClarity Administrator offers comprehensive hardware management tools that help to increase uptime, reduce costs and improve productivity through advanced server management capabilities.
- Lenovo XClarity Provisioning Manager, based in UEFI and accessible from F1 during boot, provides system inventory information, graphical UEFI Setup, platform update function, RAID Setup wizard, operating system installation function, and diagnostic functions.
- Support for Lenovo XClarity Energy Manager which captures real-time power and temperature data from the server and provides automated controls to lower energy costs.

Source: <https://lenovopress.lenovo.com/lp1396-thinksystem-sd650-n-v2-server>

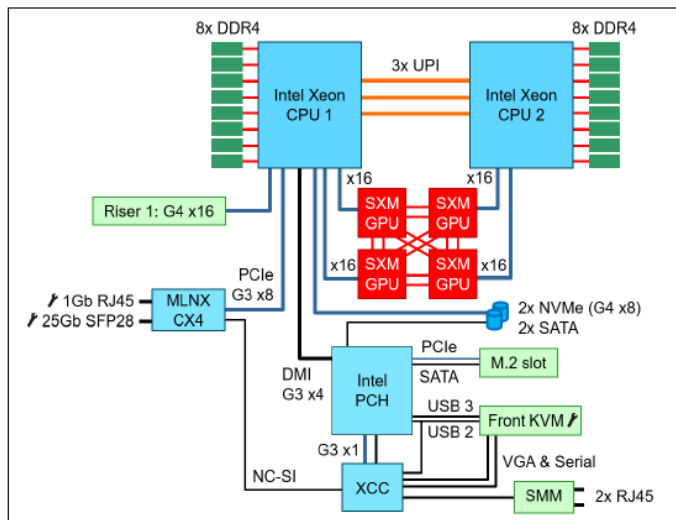
The rear of the enclosure contains the power supplies, cooling water manifolds, and the System Management Module, as shown in the following figure.



Source: <https://lenovopress.lenovo.com/lp1396-thinksystem-sd650-n-v2-server>

System architecture

The following figure shows the architectural block diagram of the SD650-N V2 with one PCIe slot and support for two drives. The GPUs each have a PCIe 4.0 x16 connection to the processors.



Source: <https://lenovopress.lenovo.com/lp1396-thinksystem-sd650-n-v2-server>

System Management Module (SMM)	The hot-swappable System Management Module (SMM2) is the management device for the enclosure. Provides integrated systems management functions and controls the power and cooling features of the enclosure. Provides remote browser and CLI-based user interfaces for remote access via the dedicated Gigabit Ethernet port. Remote access is to both the management functions of the enclosure as well as the XClarity Controller (XCC) in each server. The SMM has two Ethernet ports which enables a single incoming Ethernet connection to be daisy chained across 6 enclosures and 36 servers, thereby significantly reducing the number of Ethernet switch ports needed to manage an entire rack of SD650-N V2 servers and enclosures.
Ports	Two RJ45 port on the rear of the enclosure for 10/100/1000 Ethernet connectivity to the SMM for power and cooling management.
Systems management	Browser-based enclosure management through an Ethernet port on the SMM at the rear of the enclosure. Integrated Ethernet switch provides direct access to the XClarity Controller (XCC) embedded management of the installed servers. Servers provide more management features.

Source: <https://lenovopress.lenovo.com/lp1396-thinksystem-sd650-n-v2-server>


Lenovo ThinkSystem SN550 V2 Server

Manageability and security

The following powerful systems management features simplify the local and remote management of the SN550 V2:

- Support for Lenovo XClarity Administrator, providing auto-discovery, inventory tracking, monitoring, policy-based firmware updates, address pool management, configuration patterns and operating system installation.
- The server includes an XClarity Controller (XCC) management processor to monitor server availability and perform remote management. XCC Enterprise is supported as standard, which enables remote KVM, mounting of remote media files (ISO and IMG image files), boot capture, and power capping.
- UEFI-based Lenovo XClarity Provisioning Manager, accessible from F1 during boot, provides system inventory information, graphical UEFI Setup, platform update function, RAID Setup

Source: <https://lenovopress.lenovo.com/lp1397-thinksystem-sn550-v2-server>



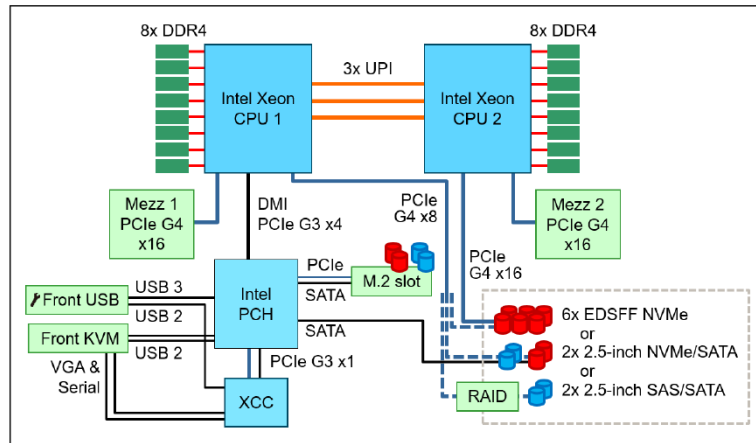
(/assets/images/LP1397/SN550%20V2%20front%20view%20with%20callouts.png)

Figure 2. Front view of the ThinkSystem SN550 V2 server

Source: <https://lenovopress.lenovo.com/lp1397-thinksystem-sn550-v2-server>

System architecture

The following figure shows the architectural block diagram of the SN550 V2, showing the major components and their connections.



(/assets/images/LP1397/SN550%20V2%20block%20diagram.png)

Figure 5. SN550 V2 system architectural block diagram

Source: <https://lenovopress.lenovo.com/lp1397-thinksystem-sn550-v2-server>

System management

The server contains an integrated service processor, XClarity Controller (XCC), which provides advanced service-processor control, monitoring, and alerting functions. The XCC is based on the Pilot4 XE401 baseboard management controller (BMC) using a dual-core ARM Cortex A9 service processor.

Local management

As shown in Figure 2, the SN550 V2 front panel includes a USB port, status indicators, a button to enable management via the USB port and a console breakout cable port. The breakout cable is supplied with the chassis and provides serial, video and two USB 2.0 ports for connecting a local console. The USB ports on the breakout cable support keyboard and mouse; storage devices are not supported.

Source: <https://lenovopress.lenovo.com/lp1397-thinksystem-sn550-v2-server>

Lenovo ThinkSystem SN550 Server (Xeon SP Gen 1)

Manageability and security

The following powerful systems management features simplify the local and remote management of the SN550:

- Support for Lenovo XClarity Administrator, providing auto-discovery, inventory tracking, monitoring, policy-based firmware updates, address pool management, configuration patterns and operating system installation.
- The server includes an XClarity Controller (XCC) to monitor server availability and perform remote management. XCC Enterprise is supported as standard, which enables remote KVM, mounting of remote media files (ISO and IMG image files), boot capture, and power capping.
- UEFI-based Lenovo XClarity Provisioning Manager, accessible from F1 during boot, provides system inventory information, graphical UEFI Setup, platform update function, RAID Setup wizard, operating system installation function, and diagnostic functions

Source: <https://lenovopress.com/lp0637-thinksystem-sn550-server-xeon-sp-gen-1>




Figure 2. Front view of the ThinkSystem SN550 Compute Node

Source: <https://lenovopress.com/lp0637-thinksystem-sn550-server-xeon-sp-gen-1>

System architecture

The following figure shows the architectural block diagram of the SN550, showing the major components and their connections.

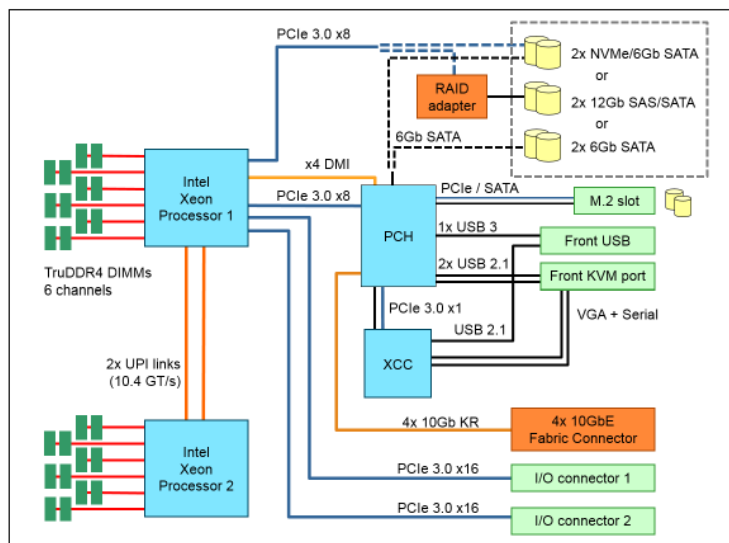


Figure 4. SN550 system architectural block diagram

Source: <https://lenovopress.com/lp0637-thinksystem-sn550-server-xeon-sp-gen-1>

System Management

The server contains an integrated service processor, XClarity Controller (XCC), which provides advanced service-processor control, monitoring, and alerting functions. The XCC is based on the Pilot4 XE401 baseboard management controller (BMC) using a dual-core ARM Cortex A9 service processor.

Local management

As shown in Figure 2, the SN550 front panel includes a USB port, status indicators, a button to enable management via the USB port and a console breakout cable port. The breakout cable supplied with the chassis provides serial, video and a USB port for connecting a local console. The USB ports on the breakout cable support keyboard and mouse; storage devices are not supported.

Source: <https://lenovopress.com/lp0637-thinksystem-sn550-server-xeon-sp-gen-1>

Lenovo XClarity Administrator™

Centralized resource-management solution designed to reduce complexity, speed response, and enhance the availability of Lenovo ThinkSystem and ThinkAgile solutions.



Lenovo XClarity Orchestrator

Lenovo XClarity Orchestrator provides centralized monitoring, management, and analytics for environments with large numbers of devices. It leverages existing XClarity Administrator across multiple instances to view overall health, collect device inventory and health summaries, drill down into device details, and view event and audit logs.



Source: <https://www.lenovo.com/us/en/data-center/software/management/>

[Lenovo XClarity Administrator > Managing servers](#)

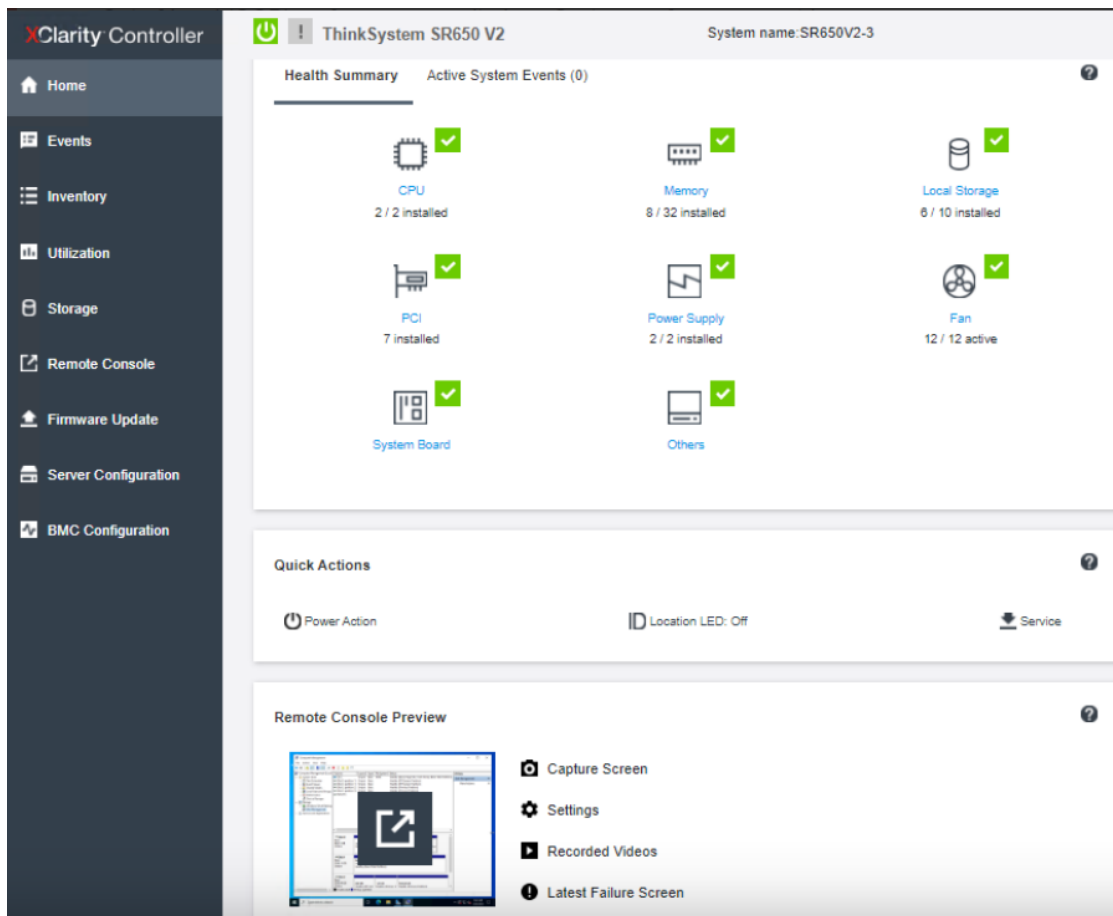
Language:

Using remote control

From the Lenovo XClarity Administrator web interface, you can open a remote-control session to a managed server as if you were at a local console. You can use the remote-control session to perform operations such as powering on or off the server, and logically mounting a local or remote drive.

Source:

https://sysmgt.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.lxca.doc%2Fremotecontrol_thinksystem_use.html/



Source:

https://sysmgt.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.lxca.doc%2Fremotecontrol_thinksystem_use.html/

Note: The XClarity Administrator does not modify the security settings or cryptographic settings (cryptographic mode and the mode used for secure communications) during the management process. You can modify the cryptographic settings after the system is managed (see [Configuring cryptography settings](#)).

Source:

https://sysmgt.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.lxca.doc%2Fsetup_manage_systems.html

6.3 Lenovo XClarity

Systems management of a cluster includes hardware management, Operating System, and Caffe applications management.

Hardware management uses the Lenovo XClarity Administrator, which is a centralized resource management solution that reduces complexity, speeds up response and enhances the availability of Lenovo server systems and solutions. XClarity is used to install the OS onto new worker nodes; update firmware across the cluster

Source: <https://lenovopress.lenovo.com/lp0892.pdf>

Lenovo XClarity Administrator provides a central interface to perform the following functions for all managed endpoints.

- **Hardware management**

Lenovo XClarity Administrator provides agent-free hardware management. It can automatically discover manageable endpoints, including Flex System chassis and components, System x, NeXtScale, and ThinkServer servers, and RackSwitch switches. Inventory of the discovered endpoints is also gathered, so an at-a-glance view of the managed hardware inventory and status is possible.

- **Configuration management**

You can quickly provision and pre-provision all of your servers using a consistent configuration. Configuration settings (such as local storage, I/O adapters, boot settings, firmware, ports, and IMM and UEFI settings) are saved as a server pattern that can be applied to one or more managed servers. When the server patterns are updated, the changes are automatically deployed to the applied servers.

- **User management**

Lenovo XClarity Administrator provides a centralized authentication server to create and manage user accounts and to manage and authenticate user credentials. The authentication server is created automatically when you start the management server for the first time. The user accounts that you create for Lenovo XClarity Administrator are also used to log in to managed chassis and servers.

Source:

<https://cc.cnetcontent.com/inlinecontent/mediaserver/test/14a/7ec/14a7ec57d22b4688923397901fb26f15/original.pdf>

Lenovo XClarity Administrator is a centralized, resource-management solution that simplifies infrastructure management, speeds responses, and enhances the availability of Lenovo® server systems and solutions. It runs as a virtual appliance that automates discovery, inventory, tracking, monitoring, and provisioning for server, network, and storage hardware in a secure environment.

Source:

https://sysmgt.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.systems.management.xcc.doc%2Fproduct_page.html&cp=2_0

Implementing a secure environment

Consider the following information when you are evaluating the security requirements for your environment:


- The physical security of your environment is important; limit access to rooms and racks where systems-management hardware is kept.
- Use a software-based firewall to protect your network hardware and data from known and emerging security threats such as viruses and unauthorized access.
- Do not change the default security settings for the network switches and pass-thru modules. The manufacturing default settings for these components disable the use of unsecure protocols and enable the requirement for signed firmware updates.
- The management applications for the CMMs, baseboard management controllers, FSPs, and switches permit only signed firmware-update packages for these components to ensure that only trusted firmware is installed.
- Only the users who are authorized to update firmware components should have firmware-update authority.

- Use the various authorization levels that are available for different users in your environment. Do not allow all users to work with the same supervisor user ID.
- Ensure that your environment meets the following NIST 800-131A criteria to support secure communications:
 - Use Secure Sockets Layer (SSL) over the TLS v1.2 protocol.
 - Use SHA-256 or stronger hashing functions for digital signatures and SHA-1 or stronger hashing functions for other applications.
 - Use RSA-2048 or stronger, or use NIST approved Elliptic Curves that are 224 bits or stronger.

Source:

https://sysmgt.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.lxca.doc%2Fsecurity_implementation.html

[Lenovo XClarity Administrator > Managing servers](#)

Language: 

Using remote control

From the Lenovo XClarity Administrator web interface, you can open a remote-control session to a managed server as if you were at a local console. You can use the remote-control session to perform operations such as powering on or off the server, and logically mounting a local or remote drive.

Source:

https://sysmgt.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.lxca.doc%2Fremotecontrol_thinksystem_use.html/

Configuring the Ethernet settings

The XClarity Controller uses two network controllers. One network controller is connected to the dedicated management port and the other network controller is connected to the shared port. Each of

Source:

https://sysmgt.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.systems.management.xcc.doc%2FN1ia_c_configuringethernet.html

Controlling access to specific devices

When devices are initially managed by Lenovo XClarity Administrator, a predefined set of role groups have permission to access the devices by default. You change the role groups that can access specific managed devices. When permission is given to certain role groups, only users that are members of those role groups can see and act on those specific devices.

Source:

https://sysmgt.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.lxca.doc%2Faccesscontrol_setspecificdevices.html

Configuring DNS

Use the information in this topic to view or change XClarity Controller Domain Name System (DNS) settings.

Note: In a Flex System, DNS settings cannot be modified on the XClarity Controller. DNS settings are managed by the CMM.

Click **Network** under **BMC Configuration** to view or modify XClarity Controller DNS settings.

Source:

https://sysmgt.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.systems.management.xcc.doc%2FN1ia_c_configuringDNS.html

Introduction

Most Lenovo ThinkSystem servers contain an integrated service processor, XClarity Controller (XCC), which provides advanced service-processor control, monitoring, and alerting functions. The XCC consolidates the service processor functionality, super I/O, video controller, and remote presence capabilities into a single chip on the server system board. The XCC is based on the Pilot4 XE401 baseboard management controller (BMC) using a dual-core ARM Cortex A9 service processor.

Source: <https://lenovopress.com/lp0880-xcc-support-on-thinksystem-servers>

67. Furthermore, the Accused '140 Products comprise at least one secure management access controller connected to one or more data buses of said managed network device for the communication of device management data. For example, the Accused '140 Products include an XCC processor, which controls remote management functions and communications regarding the same.

Introduction

Most Lenovo ThinkSystem servers contain an integrated service processor, XClarity Controller (XCC), which provides advanced service-processor control, monitoring, and alerting functions. The XCC consolidates the service processor functionality, super I/O, video controller, and remote presence capabilities into a single chip on the server system board. The XCC is based on the Pilot4 XE401 baseboard management controller (BMC) using a dual-core ARM Cortex A9 service processor.

Source: <https://lenovopress.com/lp0880-xcc-support-on-thinksystem-servers>

6.3 Lenovo XClarity

Systems management of a cluster includes hardware management, Operating System, and Caffe applications management.

Hardware management uses the Lenovo XClarity Administrator, which is a centralized resource management solution that reduces complexity, speeds up response and enhances the availability of Lenovo server systems and solutions. XClarity is used to install the OS onto new worker nodes; update firmware across the cluster

Source: <https://lenovopress.lenovo.com/lp0892.pdf>

System architecture

The following figure shows the architectural block diagram of the SN550, showing the major components and their connections.

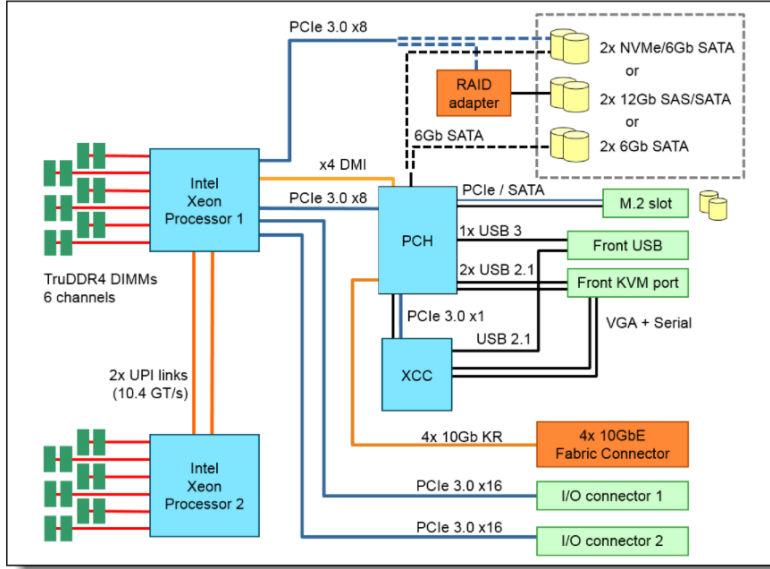
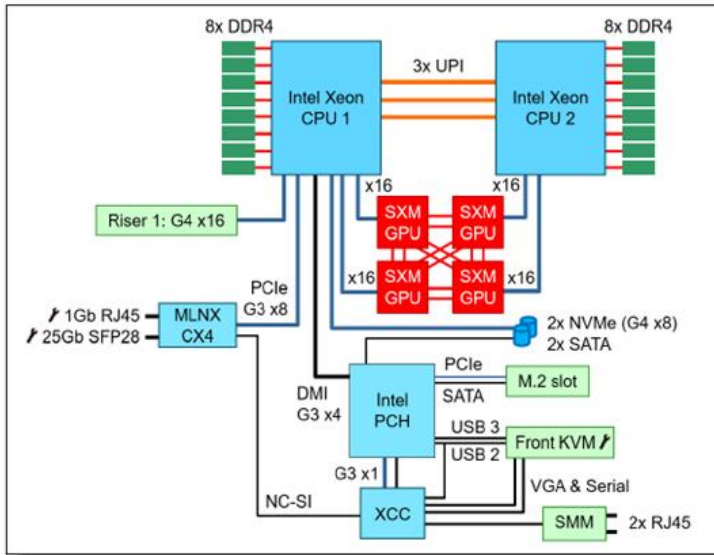


Figure 4. SN550 system architectural block diagram

Source: <https://lenovopress.com/lp0637-thinksystem-sn550-server-xeon-sp-gen-1>

System architecture

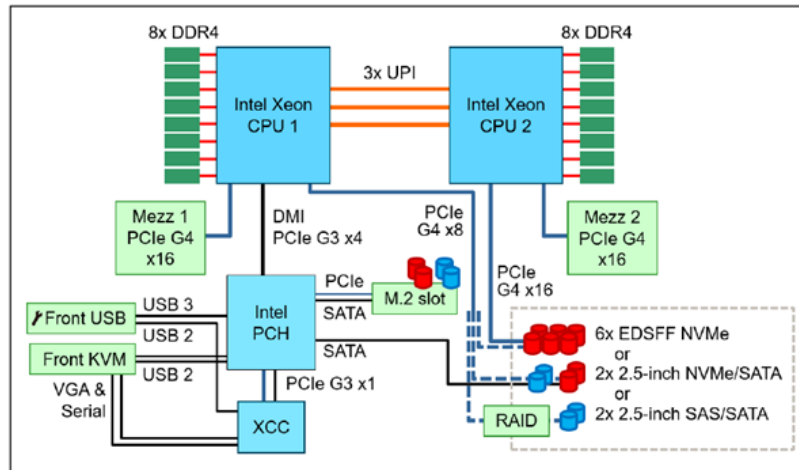
The following figure shows the architectural block diagram of the SD650-N V2 with one PCIe slot and support for two drives. The GPUs each have a PCIe 4.0 x16 connection to the processors.



Source: <https://lenovopress.lenovo.com/lp1397-thinksystem-sn550-v2-server>

System architecture

The following figure shows the architectural block diagram of the SN550 V2, showing the major components and their connections.



(/assets/images/LP1397/SN550%20V2%20block%20diagram.png)

Figure 5. SN550 V2 system architectural block diagram

Source: <https://lenovopress.lenovo.com/lp1396-thinksystem-sd650-n-v2-server>

[Lenovo XClarity Administrator](#) >
[Managing servers](#)

Language:

Using remote control

From the Lenovo XClarity Administrator web interface, you can open a remote-control session to a managed server as if you were at a local console. You can use the remote-control session to perform operations such as powering on or off the server, and logically mounting a local or remote drive.

Source:

https://sysmgt.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.lxca.doc%2Fremotecontrol_thinksystem_use.html/

Lenovo XClarity Administrator provides a central interface to perform the following functions for all managed endpoints.

- **Hardware management**

Lenovo XClarity Administrator provides agent-free hardware management. It can automatically discover manageable endpoints, including Flex System chassis and components, System x, NeXtScale, and ThinkServer servers, and RackSwitch switches. Inventory of the discovered endpoints is also gathered, so an at-a-glance view of the managed hardware inventory and status is possible.

- **Configuration management**

You can quickly provision and pre-provision all of your servers using a consistent configuration. Configuration settings (such as local storage, I/O adapters, boot settings, firmware, ports, and IMM and UEFI settings) are saved as a server pattern that can be applied to one or more managed servers. When the server patterns are updated, the changes are automatically deployed to the applied servers.

- **User management**

Lenovo XClarity Administrator provides a centralized authentication server to create and manage user accounts and to manage and authenticate user credentials. The authentication server is created automatically when you start the management server for the first time. The user accounts that you create for Lenovo XClarity Administrator are also used to log in to managed chassis and servers.

Source:

<https://cc.cnetcontent.com/inlinecontent/mediaserver/test/14a/7ec/14a7ec57d22b4688923397901fb26f15/original.pdf>

68. The Accused '140 Products further comprise an out-of-band access connection means for connecting said one or more network services or remote users with said secure management access controller for management of said network device. For example, connectivity to the XCC processor for remote device management can be over an out-of-band management connection, giving remote administrators a secure out-of-band management solution, as illustrated below:

6.3 Lenovo XClarity

Systems management of a cluster includes hardware management, Operating System, and Caffe applications management.

Hardware management uses the Lenovo XClarity Administrator, which is a centralized resource management solution that reduces complexity, speeds up response and enhances the availability of Lenovo server systems and solutions. XClarity is used to install the OS onto new worker nodes; update firmware across the cluster

Source: <https://lenovopress.lenovo.com/lp0892.pdf>

Updating a storage HBA controller's firmware through out-of-band(OOB)channels like XCC Web GUI might fail.

(where HBA = Host Bus Adapter, XCC = Lenovo XClarity Controller)

When system POST completes, XCC starts to scan each storage controller in the system. If a storage HBA controller does not respond to XCC correctly at that time, XCC will mark it as problematic. Afterwards, if user try to update the controller's firmware through an OOB channel, a failure will occur.

Source:

<https://datacentersupport.lenovo.com/il/en/products/servers/thinksystem/sr650/7x05/solutions/ht508305>

Introduction

Most Lenovo ThinkSystem servers contain an integrated service processor, XClarity Controller (XCC), which provides advanced service-processor control, monitoring, and alerting functions. The XCC consolidates the service processor functionality, super I/O, video controller, and remote presence capabilities into a single chip on the server system board. The XCC is based on the Pilot4 XE401 baseboard management controller (BMC) using a dual-core ARM Cortex A9 service processor.

Source: <https://lenovopress.com/lp0880-xcc-support-on-thinksystem-servers>

Controlling access to specific devices

When devices are initially managed by Lenovo XClarity Administrator, a predefined set of role groups have permission to access the devices by default. You change the role groups that can access specific managed devices. When permission is given to certain role groups, only users that are members of those role groups can see and act on those specific devices.

Source:

https://sysmgmt.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.lxca.doc%2Faccesscontrol_setspecificdevices.html

Isolated management network

The use of a separate management network is always a preferred practice for the isolation of data and management environments. An isolated management network can be used in a lights-out environment to provide out-of-band connectivity to locate and troubleshoot issues that might span across the data network and multiple servers. In today's data centers, this network often consists of a dedicated management switch that uses 1 Gb connectivity.

Figure 7 shows Lenovo Switching products that are connecting to a separate 1 Gb Management for out-of-band management. Server management ports (XCC and IMM2), as well as switch management ports, should be connected to this out-of-band network so that they can be reachable in the event of an outage or another issue on the main data network.




Figure 7 Out of Band 1G Management connectivity

Source: <https://lenovopress.lenovo.com/lp1068.pdf>

Configuring a dedicated or shared network port

The SR950 offers two RJ45 ports that you can use to access XCC remotely, either the dedicated management port shown in Figure 5-27 on page 143, or via a port on an installed Ethernet LOM adapter which will be shared with the installed operating system. You select the port by setting Dedicated or Shared for the Network Interface Port in the Network Configuration panel of F1 setup.

Source: <https://lenovopress.lenovo.com/lp0746.pdf>

Default local network access to XCC

The default network connection for the XCC on the SR950 is through the System Management port on the back of the server as shown in Figure 5-27.




Figure 5-27 XClarity Controller (XCC) port on the rear of the SR950

The XCC network access label is on the front of the server accessible via a pull-out tab. The label provides the default IPv4 host name and default IPv6 link local address of the XCC. A sample label is shown in Figure 5-28 on page 144.

Source: <https://lenovopress.lenovo.com/lp0746.pdf>

69. In addition, the Accused '140 Products include at least one virtual management interface connection means for connecting said one or more network services or remote users with said secure management controller. For example, connectivity to the XCC processor for remote device management can be over shared network connections via a shared network port, giving remote administrators a secure in-band management solution that virtually separates user and management traffic. Furthermore, in the Accused '140 Products, the virtual management interface connection means provides logical separation of management data from user data and utilizes user interfaces of said managed network element for connecting said one or more network services or remote users with said secure management access controller. For example, the shared network port allows remote administrators to communicate management data and logically separate the user data from the management data via virtualization and VLAN tagging, as seen below:

Management interfaces

There are two ways to access the XCC management processor remotely:

- Command-line interface. To access the CLI interface, use SSH to log in to the management processor.
- Web-based interface. To access the web-based interface, point your browser to the IP address for the management processor. The new intuitive interface includes at-a-glance visualizations and simple access to common system actions. The dashboard is shown in the following figure.

Source: <https://lenovopress.lenovo.com/lp0880-xcc-support-on-thinksystem-servers>

Configuring a dedicated or shared network port

The SR950 offers two RJ45 ports that you can use to access XCC remotely, either the dedicated management port shown in Figure 5-27 on page 143, or via a port on an installed Ethernet LOM adapter which will be shared with the installed operating system. You select the port by setting Dedicated or Shared for the Network Interface Port in the Network Configuration panel of F1 setup.

Refer to 5.1.1, “Accessing Lenovo XClarity Provisioning Manager” on page 120 for information on how to access the F1 setup options (LXPM). Once in LXPM, you can access this panel by selecting **UEFI Setup** → **BMC Settings** → **Network Settings**, as shown in Figure 5-29.




Figure 5-29 Configuring dedicated or shared XCC port

Source: <https://lenovopress.lenovo.com/lp0746.pdf>

XCC dedicated versus shared Ethernet port

When configured as **Dedicated**, you are connecting to the network via the system management port as shown in Figure 5-27 on page 143.

The use of this port allows for easier separation of public and management network traffic. Separating the traffic is done when you connect your public network port to switch ports that belong to a public access virtual LAN (VLAN). The management port is connected to a switch port defined by a separate management VLAN.

When configured as **Shared**, you are sharing network traffic between the management port and on an Ethernet adapter.

Although the Shared configuration eliminates a physical switch port and patch cable configuration, the media access control (MAC) address for the shared Ethernet port and the MAC address for the XCC address through this single network port. This situation means that there are at least two separate IP addresses for the same physical port, which prevents you from configuring the other adapter's Ethernet port in a network team by using 802.3ad load balancing.

To maintain separation between public and management traffic, network teaming software must be used to establish a VLAN to be used by the server to send public-tagged traffic to the network switch. The switch port must be configured as a trunk port to support the public-tagged VLAN traffic and the untagged traffic for the management. The management VLAN must be defined as the native VLAN on the switch port so that its untagged traffic from the switch is accepted by the XCC MAC and dropped by the second Ethernet port's MAC.

Source: <https://lenovopress.lenovo.com/lp0746.pdf>

Configuring the Ethernet settings

Use the information in this topic to view or change how the XClarity Controller communicates by way of an Ethernet connection.

The XClarity Controller uses two network controllers. One network controller is connected to the dedicated management port and the other network controller is connected to the shared port. Each of the network controllers is assigned its own burned in MAC address. If DHCP is being used to assign an IP address to the XClarity Controller, when a user switches between network ports or when a failover from the dedicated network port to the shared network port occurs, a different IP address may be assigned to the XClarity Controller by the DHCP server. It is recommended that when using DHCP, users should use the host name to access the XClarity Controller rather than relying on an IP address. Even if the XClarity Controller network ports are not changed, the DHCP server could possibly assign a different IP address to the XClarity Controller when the DHCP lease expires, or when the XClarity Controller reboots. If a user needs to access the XClarity Controller using an IP address that will not change, the XClarity Controller should be configured for a static IP address rather than DHCP.

Source:

https://sysmgt.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.systems.management.xcc.doc%2FN1ia_c_configuringethernet.html

After configuring the BMC settings as shown in the screenshot below, XCC access will be activated on the shared NIC if the connection through the XCC dedicated port fails, but the XCC dedicated port will no longer be available after recovering the connection or the connection through the shared NIC fails.

Source: <https://support.lenovo.com/in/en/solutions/ht510765-xcc-cannot-be-accessed-from-dedicated-port-after-failover-to-the-shared-nic-lenovo-thinksystem>

Configuring advanced Ethernet settings

Click the **Advanced Ethernet** tab to set additional Ethernet settings.

Note: In a Flex System, the VLAN settings are managed by a Flex System CMM and cannot be modified on the XClarity Controller.

To enable Virtual LAN (VLAN) tagging select the **Enable VLAN** check box. When VLAN is enabled and a VLAN ID is configured, the XClarity Controller only accepts packets with the specified VLAN IDs. The VLAN IDs can be configured with numeric values between 1 and 4094.

Source:

https://sysmgt.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.systems.management.xcc.doc%2FNFIa_c_configuringethernet.html

70. Additionally, Defendant has been, and currently is, an active inducer of infringement of the '140 patent under 35 U.S.C. § 271(b) and a contributory infringer the '140 patent under 35 U.S.C. § 271(c).

71. Defendant has actively induced, and continues to actively induce, infringement of the '140 patent by causing others to use, offer for sale, or sell in the United States, products or services covered by the '140 patent, including but not limited to the '140 Accused Products and any other products or services that include XCC, and/or IMM2 processors, or processors with the functionality described above. Defendant provides these products and services to others, such as customers, resellers, partners, and end-users, who, in turn, use, provision for use, offer for sale, or sell those products and services, which directly infringe the '140 patent as described above. Defendant's inducement includes the directions and instructions found at one or more of the following links, the provision of which was on-going as of the filing of the Complaint in case 6:23-cv-0068-ADA, and remains on-going as of the filing of this Complaint, and the content of which is specifically illustrated above:

- <https://lenovopress.com/lp0637-thinksystem-sn550-server-xeon-sp-gen-1>
- <https://lenovopress.com/lp0880-xcc-support-on-thinksystem-servers>
- https://sysmgt.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.lxca.doc%2Fremot_econtrol_thinksystem_use.html
- https://sysmgt.lenovofiles.com/help/index.jsp?topic=%2Flxcc_frontend%2Flxcc_overview.html&cp=3

- https://sysmgt.lenovofiles.com/help/topic/com.lenovo.systems.management.xcc.doc/NN1ia_c_configuringethernet.html?cp=3_0_3_1_0
- https://sysmgt.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.systems.management.xcc.doc%2Fdw1lm_c_accessingtheimmwebinterface.html&cp=3_0_2_0
- https://sysmgt.lenovofiles.com/help/topic/com.lenovo.systems.management.xcc.doc/dw1lm_t_settinguptheimmnetworkconnection.html?cp=3_0_2_0_0
- https://sysmgt.lenovofiles.com/help/topic/com.lenovo.systems.management.xcc.doc/NN1ia_c_configuringethernet.html
- https://sysmgt.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.lxca.doc%2Fsecurity_implement.html
- https://sysmgt.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.lxca.doc%2Faccess_control_setspecificdevices.html
- https://sysmgt.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.systems.management.xcc.doc%2FNN1ia_c_configuringDNS.html
- https://sysmgt.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.lxca.doc%2Fsecurity_implement.html
- https://sysmgt.lenovofiles.com/help/topic/com.lenovo.lxca.doc/lxca_qsg_install_kvm.pdf
- <https://www.lenovo.com/in/en/data-center/software/systems-management/c/systems-management>
- <https://www.lenovo.com/us/en/data-center/software/management/>
- https://systemx.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.systemx.common.nav.doc%2Foverview_imm2.html&cp=0_4
- https://systemx.lenovofiles.com/help/topic/com.lenovo.sysx.imm2.doc/dw1lm_c_ch3_configuringtheimm.html
- https://systemx.lenovofiles.com/help/topic/com.lenovo.sysx.imm2.doc/NN1ia_c_configuringnetworkprotocolproperties.html
- <https://lenovopress.lenovo.com/lp0746.pdf>

72. Defendant has contributed to, and continues to contribute to, the infringement of the '140 patent by others by knowingly providing one or more components, for example the XCC and/or IMM2 processors included in the Accused Products, a portion thereof, and/or the software/hardware modules responsible for the accused functionality described herein, that, when installed, configured, and used, result in systems that, as intended by Defendant described above, directly infringe one or more claims of the '140 patent.

73. Defendant knew of the '140 patent, or should have known of the '140 patent, but was willfully blind to its existence. Upon information and belief, Defendant had actual knowledge of the '140 patent at least as early as February 2, 2023, the date the Complaint in case 6:23-cv-0068-ADA was filed or as early as February 3, 2023, the date upon which IV notified Defendant of the filing of the Complaint in the aforementioned action. Alternatively, upon information and belief, Defendant has had actual knowledge of the '140 patent since receipt of this Complaint or service on Defendant of this Complaint.

74. By the time of trial, Defendant will or should have known and intended (since receiving such notice) that its continued actions would infringe and would actively induce and contribute to the infringement of the '140 patent.

75. Defendant has committed, and continues to commit, contributory infringement by selling products and services that directly infringe the '140 patent when used by a third party, such as the Accused '140 Products, and that are a material part of the invention, knowing them to be especially made or adapted for use in infringement of the '140 patent and not staple articles or commodities of commerce suitable for substantial non-infringing use.

76. As a result of Defendant's acts of infringement, IV has suffered and will continue to suffer damages in an amount to be determined at trial.

COUNT II

(Defendant's Infringement of U.S. Patent No. 8,474,016)

77. The preceding paragraphs are reincorporated by reference as if fully set forth herein.

78. The '016 patent claims and teaches, *inter alia*, an improved apparatus and method to provide secure remote management for devices by deeply embedding the necessary secure remote management hardware and software in the managed device itself. In so doing, the

apparatus and methods separate user traffic from device management traffic, logically and physically, both in the device and while in transit over a network.

79. The inventions improved upon then-existing remote access/management security techniques by providing for the separation of management data from user data both in the device being managed and while the management data is in transit. There was, at the time, a recognized need for securing remote management and for separating management data from user data; yet this had not been done before the inventions disclosed in the '016.

80. The '016 patent solved this long-felt need, among others, by combining in a network device a bus controller that is connected to two separate buses, wherein the bus controller receives encrypted management requests from the second bus and sends them to the first bus, and a processor that decrypts management requests from the first bus, wherein the claimed processor is separate from a processor on the network device, all the while without requiring additional devices taking up additional rack space by embedding the necessary hardware and software for secure management of the device in the device to be managed.

81. More specifically, the claims of the '016 patent recite a processor configured to control one or more functions of a network device having a network interface, wherein the network device is configured to receive data requests and an encrypted form of management requests via the network interface, wherein the management requests are from a remote administrator; a first bus; a bus controller coupled to the processor via the first bus, wherein the bus controller is also coupled to a second bus of the network device that is distinct from the first bus; wherein the bus controller is configured to receive the encrypted form of the management requests from the second bus, and to convey the encrypted form of the management requests to the processor via the first bus; wherein the processor is configured to decrypt the encrypted form of the management

requests; wherein the network device includes a processor configured to facilitate operation of the network device; and wherein the processor of the apparatus is distinct from the processor included in the network device.

82. The apparatus covered by the asserted claims, therefore, differs markedly from the prior systems in use at the time of this invention, which lacked the claimed combination of a processor configured to control one or more functions of a network device having a network interface, wherein the network device is configured to receive data requests and an encrypted form of management requests via the network interface, wherein the management requests are from a remote administrator; a first bus; a bus controller coupled to the processor via the first bus, wherein the bus controller is also coupled to a second bus of the network device that is distinct from the first bus; wherein the bus controller is configured to receive the encrypted form of the management requests from the second bus, and to convey the encrypted form of the management requests to the processor via the first bus; wherein the processor is configured to decrypt the encrypted form of the management requests; wherein the network device includes a processor configured to facilitate operation of the network device; and wherein the processor of the apparatus is distinct from the processor included in the network device.

83. Defendant has directly infringed and continues to directly infringe at least claim 1 of the '016 patent by making, using, testing, selling, offering for sale, and importing into the United States products and services covered by one or more claims of the '016 patent. Defendant's products and services that infringe the '016 patent include all products and services that use an XCC embedded processor, which include the ThinkSystem server series, and any chassis/enclosures such products may be housed in, as well as any other of Defendant's products

and services, either alone or in combination, that operate in substantially the same manner (together, the “Accused ’016 Products” or “Accused Products”).

84. Claim 1 of the ’016 patent is reproduced below:

1. An apparatus, comprising:

a processor configured to control one or more functions of a network device having a network interface, wherein the network device is configured to receive data requests and an encrypted form of management requests via the network interface, wherein the management requests are from a remote administrator;

a first bus; and

a bus controller coupled to the processor via the first bus, wherein the bus controller is also coupled to a second bus of the network device that is distinct from the first bus, wherein the bus controller is configured to receive the encrypted form of the management requests from the second bus, and to convey the encrypted form of the management requests to the processor via the first bus;

wherein the processor is configured to decrypt the encrypted form of the management requests, wherein the network device includes a processor configured to facilitate operation of the network device, and wherein the processor of the apparatus is distinct from the processor included in the network device.

85. The Accused ’016 Products each provide a processor configured to control one or more functions of a network device having a network interface. For example, the Accused ’016 Products are network devices, modules, and nodes capable of being configured in a distributed computer network, such as the Lenovo ThinkSystem SN550 Server (Xeon SP Gen 2), SN550 V2 Server and SD650-N V2 Server that support the XCC Standard, XCC Advanced and XCC Enterprise integrated service processors for controlling and securing remote management applications and services as well as communications regarding the same, as seen below:

Lenovo XClarity Controller

Language:

The Lenovo XClarity Controller (XCC) is the next generation management controller for Lenovo ThinkSystem servers.

- [Lenovo XClarity Controller with Intel Xeon SP \(1st, 2nd Gen\)](#).
- [Lenovo XClarity Controller with Intel Xeon SP \(3rd Gen\) and AMD EPYC \(2nd, 3rd Gen\)](#).
- [Lenovo XClarity Controller REST API reference](#)

The following table lists the server models supported by LXCC Products:

Table 1. Server models supported by LXCC

Lenovo XClarity Controller	Server models	
Lenovo XClarity Controller with Intel Xeon SP (1st, 2nd Gen)	<ul style="list-style-type: none"> • SD530 • SD650 • SE350 • SN550 • SN850 • ST250 • ST258 • ST550 • ST558 	<ul style="list-style-type: none"> • SR570 • SR590 • SR630 • SR650 • SR670 • SR850 • SR850p • SR860 • SR950
Lenovo XClarity Controller with Intel Xeon SP (3rd Gen) and AMD EPYC (2nd, 3rd Gen)	<ul style="list-style-type: none"> • SD630 V2 • SD650 V2 • SD650-N V2 • SN550 V2 • ST650 V2 	<ul style="list-style-type: none"> • SR630 V2 • SR645 • SR650 V2 • SR665 • SR670 V2 • SR850 V2 • SR860 V2

Source:

https://sysmgmt.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.lxca.doc%2Fremotecontrol_thinksystem_use.html

Server support

The following table shows what level of XCC is included with each ThinkSystem server.

Table 3. Server support

Server	XCC Standard	XCC Advanced	XCC Enterprise
ThinkSystem V2 servers			
ST650 V2 (7Z74/7Z75)	Supported	Supported	Supported
SR630 V2 (7Z70/7Z71)	Supported	Supported	Supported
SR650 V2 (7Z72/7Z73)	Supported	Supported	Supported
SR670 V2 (7Z22/7Z23)	Supported	Upgrade	Upgrade
SR850 V2 (7D31 / 7D32 / 7D33)	N/A	N/A	All models
SR860 V2 (7Z59 / 7Z60)	N/A	N/A	All models
SD630 V2 (7D1K)	Supported	Upgrade	Upgrade
SD650 V2 (7D1M)	Supported	Upgrade	Upgrade
SD650-N V2 (7D1N)	Supported	Upgrade	Upgrade
SN550 V2 (7Z69)	N/A	N/A	All models
ThinkSystem V1 servers			
SE350 (7Z46 / 7D1X)	Supported	Most models**	Some models**
ST50 (7Y48/7Y50)	Not supported	Not supported	Not supported
ST250 (7Y45/7Y46)	Most models*	Upgrade	Upgrade
SR150 (7Y54)	Most models*	Upgrade	Upgrade
SR250 (7Y51/7Y52)	Most models*	Upgrade	Upgrade
ST550 (7X09 / 7X10)	Most models*	Upgrade	Upgrade
SR530 (7X07 / 7X08)	Most models*	Upgrade	Upgrade
SR550 (7X03 / 7X04)	Most models*	Upgrade	Upgrade
SR570 (7Y02 / 7Y03)	Most models*	Upgrade	Upgrade
SR590 (7X98 / 7X99)	Most models*	Upgrade	Upgrade
SR630 (7X01 / 7X02)	Most models*	Upgrade	Upgrade
SR635 (7Y98 / 7Y99)	Not supported	Not supported	Not supported
SR645 (7D2Y/7D2X)	Most models*	Upgrade	Upgrade
SR650 (7X05 / 7X06)	Most models*	Upgrade	Upgrade
SR655 (7Y00 / 7Z01)	Not supported	Not supported	Not supported
SR665 (7D2W/7D2V)	Most models*	Upgrade	Upgrade
SR670 (7Y36 / 7Y37 / 7Y38)	Configure-to-order	Configure-to-order	Configure-to-order
SR850 (7X18 / 7X19)	N/A	Most models**	Upgrade
SR850P (7D2F / 2D2G)	N/A	N/A	Most models**
SR860 (7X69 / 7X70)	N/A	Most models**	Upgrade
SR950 (7X11 / 7X12 / 7X13)	N/A	N/A	All models
SD530 (7X21)	Most models*	Upgrade	Upgrade
SD650 (7X58)	Configure-to-order	Configure-to-order	Configure-to-order
SN550 (7X16)	N/A	N/A	All models
SN850 (7X15)	N/A	N/A	All models

Part numbers

Models of ThinkSystem servers come with either XClarity Controller Standard, Advanced or Enterprise, depending on the server type and the model. The servers will be delivered with the stated version already active. The following table shows the field upgrades available for models that come with XCC Standard or XCC Advanced.

Source: <https://lenovopress.com/lp0880-xcc-support-on-thinksystem-servers>



Lenovo ThinkSystem SN550 Server (Xeon SP Gen 2) Product Guide

Manageability and security

The following powerful systems management features simplify the local and remote management of the SN550:

- Support for Lenovo XClarity Administrator, providing auto-discovery, inventory tracking, monitoring, policy-based firmware updates, address pool management, configuration patterns and operating system installation.
- The server includes an XClarity Controller (XCC) management processor to monitor server availability and perform remote management. XCC Enterprise is supported as standard, which enables remote KVM, mounting of remote media files (ISO and IMG image files), boot capture, and power capping.
- UEFI-based Lenovo XClarity Provisioning Manager, accessible from F1 during boot, provides system inventory information, graphical UEFI Setup, platform update function, RAID Setup wizard, operating system installation function, and diagnostic functions

Source: <https://lenovopress.lenovo.com/lp1043.pdf>




Figure 2. Front view of the ThinkSystem SN550 Compute Node

Source: <https://lenovopress.lenovo.com/lp1043.pdf>

System architecture

The following figure shows the architectural block diagram of the SN550, showing the major components and their connections.

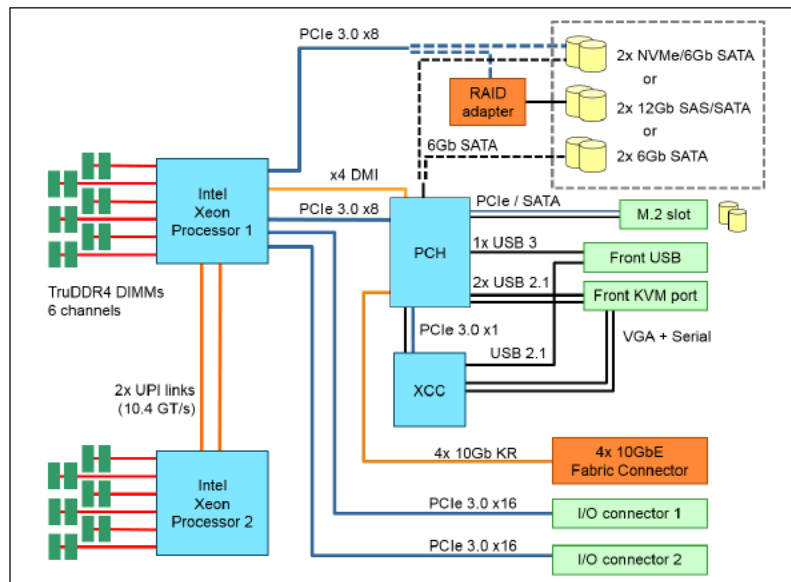


Figure 4. SN550 system architectural block diagram

Source: <https://lenovopress.lenovo.com/lp1043.pdf>

System Management

The server contains an integrated service processor, XClarity Controller (XCC), which provides advanced service-processor control, monitoring, and alerting functions. The XCC is based on the Pilot4 XE401 baseboard management controller (BMC) using a dual-core ARM Cortex A9 service processor.

Local management

As shown in Figure 2, the SN550 front panel includes a USB port, status indicators, a button to enable management via the USB port and a console breakout cable port. The breakout cable supplied with the chassis provides serial, video and a USB port for connecting a local console. The USB ports on the breakout cable support keyboard and mouse; storage devices are not supported.

Source: <https://lenovopress.lenovo.com/lp1043.pdf>


Lenovo ThinkSystem SN550 V2 Server

Manageability and security

The following powerful systems management features simplify the local and remote management of the SN550 V2:

- Support for Lenovo XClarity Administrator, providing auto-discovery, inventory tracking, monitoring, policy-based firmware updates, address pool management, configuration patterns and operating system installation.
- The server includes an XClarity Controller (XCC) management processor to monitor server availability and perform remote management. XCC Enterprise is supported as standard, which enables remote KVM, mounting of remote media files (ISO and IMG image files), boot capture, and power capping.
- UEFI-based Lenovo XClarity Provisioning Manager, accessible from F1 during boot, provides system inventory information, graphical UEFI Setup, platform update function, RAID Setup

Source: <https://lenovopress.lenovo.com/lp1397-thinksystem-sn550-v2-server>



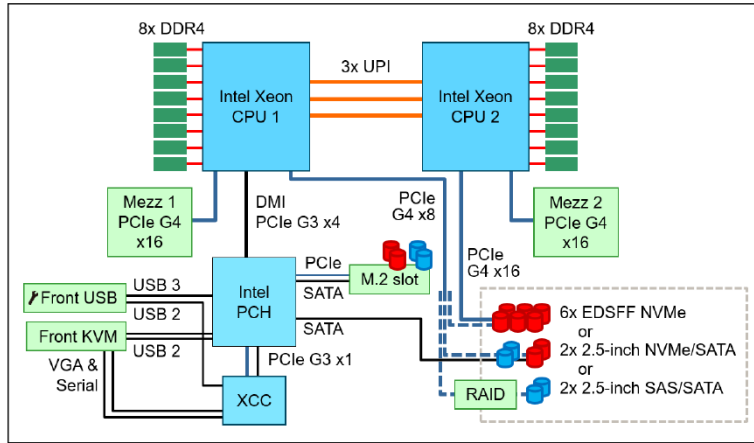
(/assets/images/LP1397/SN550%20V2%20front%20view%20with%20callouts.png)

Figure 2. Front view of the ThinkSystem SN550 V2 server

Source: <https://lenovopress.lenovo.com/lp1397-thinksystem-sn550-v2-server>

System architecture

The following figure shows the architectural block diagram of the SN550 V2, showing the major components and their connections.



(/assets/images/LP1397/SN550%20V2%20block%20diagram.png)

Figure 5. SN550 V2 system architectural block diagram

Source: <https://lenovopress.lenovo.com/lp1397-thinksystem-sn550-v2-server>

System management

The server contains an integrated service processor, XClarity Controller (XCC), which provides advanced service-processor control, monitoring, and alerting functions. The XCC is based on the Pilot4 XE401 baseboard management controller (BMC) using a dual-core ARM Cortex A9 service processor.

Source: <https://lenovopress.lenovo.com/lp1397-thinksystem-sn550-v2-server>


Lenovo ThinkSystem SD650-N V2 Server

Manageability and security

The following powerful systems management features simplify local and remote management of the SD650-N V2 server:

- The server includes an XClarity Controller (XCC) to monitor server availability. Optional upgrade to XCC Advanced to provide remote control (keyboard video mouse) functions. Optional upgrade to XCC Enterprise enables the additional support for the mounting of remote media files (ISO and IMG image files), boot capture, and power capping.
- Lenovo XClarity Administrator offers comprehensive hardware management tools that help to increase uptime, reduce costs and improve productivity through advanced server management capabilities.
- Lenovo XClarity Provisioning Manager, based in UEFI and accessible from F1 during boot, provides system inventory information, graphical UEFI Setup, platform update function, RAID Setup wizard, operating system installation function, and diagnostic functions.
- Support for Lenovo XClarity Energy Manager which captures real-time power and temperature data from the server and provides automated controls to lower energy costs.

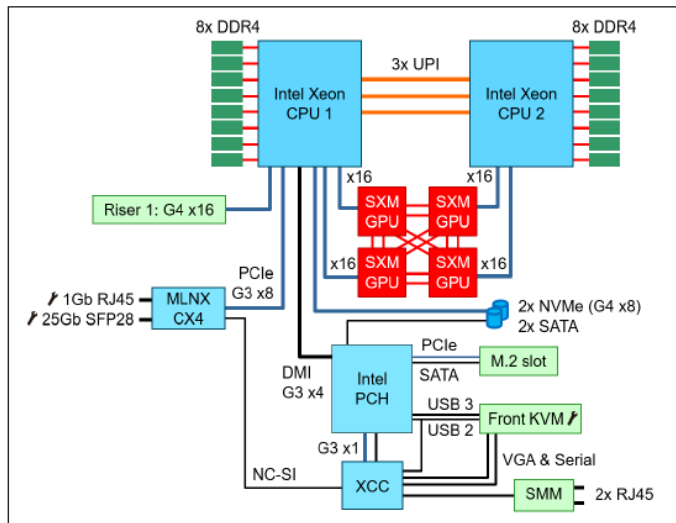
Source: <https://lenovopress.lenovo.com/lp1396-thinksystem-sd650-n-v2-server>



Source: <https://lenovopress.lenovo.com/lp1396-thinksystem-sd650-n-v2-server>

System architecture

The following figure shows the architectural block diagram of the SD650-N V2 with one PCIe slot and support for two drives. The GPUs each have a PCIe 4.0 x16 connection to the processors.



Source: <https://lenovopress.lenovo.com/lp1396-thinksystem-sd650-n-v2-server>

<p>System Management Module (SMM)</p>	<p>The hot-swappable System Management Module (SMM2) is the management device for the enclosure. Provides integrated systems management functions and controls the power and cooling features of the enclosure. Provides remote browser and CLI-based user interfaces for remote access via the dedicated Gigabit Ethernet port. Remote access is to both the management functions of the enclosure as well as the XClarity Controller (XCC) in each server.</p> <p>The SMM has two Ethernet ports which enables a single incoming Ethernet connection to be daisy chained across 6 enclosures and 36 servers, thereby significantly reducing the number of Ethernet switch ports needed to manage an entire rack of SD650-N V2 servers and enclosures.</p>
<p>Ports</p>	<p>Two RJ45 port on the rear of the enclosure for 10/100/1000 Ethernet connectivity to the SMM for power and cooling management.</p>

Systems management	Browser-based enclosure management through an Ethernet port on the SMM at the rear of the enclosure. Integrated Ethernet switch provides direct access to the XClarity Controller (XCC) embedded management of the installed servers. Servers provide more management features.
--------------------	---

Source: <https://lenovopress.lenovo.com/lp1396-thinksystem-sd650-n-v2-server>

Lenovo XClarity Controller is embedded in every ThinkSystem server on a separate microprocessor, and is designed to help you standardize, simplify and automate foundation server management tasks.

Source: <https://www.lenovo.com/in/en/data-center/software/systems-management/XClarity-Controller/p/WMD00000367?orgRef=https%253A%252F%252Fwww.google.com%252F>

The Lenovo XClarity Controller (XCC) is the next generation management controller that replaces the baseboard management controller (BMC) for Lenovo ThinkSystem servers.

It is the follow-on to the Integrated Management Module II (IMM2) service processor that consolidates the service processor functionality, Super I/O, video controller, and remote presence capabilities into a single chip on the server system board. It provides features such as the following:

- Choice of a dedicated or shared Ethernet connection for systems management
- Support for HTML5
- Support for access via XClarity Mobile
- XClarity Provisioning Manager
- Remote configuration using XClarity Essentials or XClarity Controller CLI.
- Capability for applications and tools to access the XClarity Controller either locally or remotely
- Enhanced remote-presence capabilities.
- REST API (Redfish schema) support for additional web-related services and software applications.

Source:

https://sysmgt.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.systems.management.xcc.doc%2FNFI1ia_c_configuringsecurity.html

86. In the Accused '016 Products, the network device is configured to receive data requests and an encrypted form of management requests via the network interface, wherein the management requests are from a remote administrator. For example, the ThinkSystem device receives remote management requests on the 4x 10GbE fabric connector (or equivalent network interface), which is encrypted for security purposes, as can be seen below:

System architecture

The following figure shows the architectural block diagram of the SN550, showing the major components and their connections.

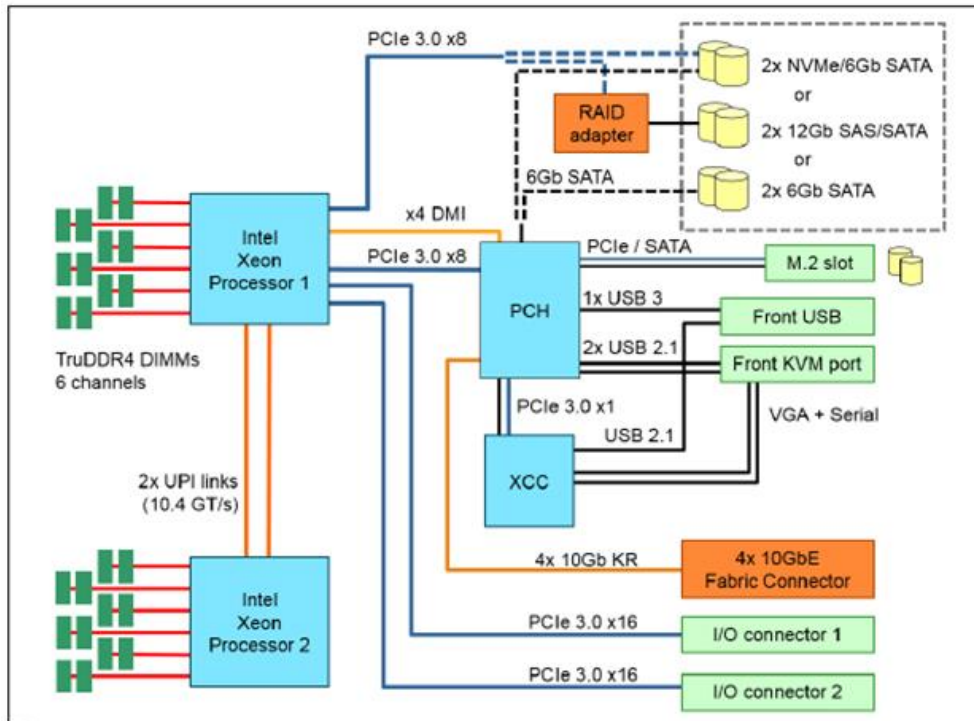


Figure 4. SN550 system architectural block diagram

Source: <https://lenovopress.lenovo.com/lp1043.pdf>

Embedded 10Gb Network Adapter

The SN550 includes an embedded 4-port 10Gb Intel controller built into the system board. As listed in the [Models](#) section, some SN550 models include the Fabric Connector needed to connect the embedded controller to the midplane of the Flex System chassis. For models that do not include the Fabric Connector, it can be ordered and installed in the field. Ordering information is listed in the following table.

Source: <https://lenovopress.lenovo.com/lp1043.pdf>

The following figure shows the location of the I/O module bays in the Flex System Enterprise Chassis.




Figure 8. Location of the I/O module bays in the Flex System Enterprise Chassis

The following figure shows how adapters are connected to I/O modules that are installed in the chassis.

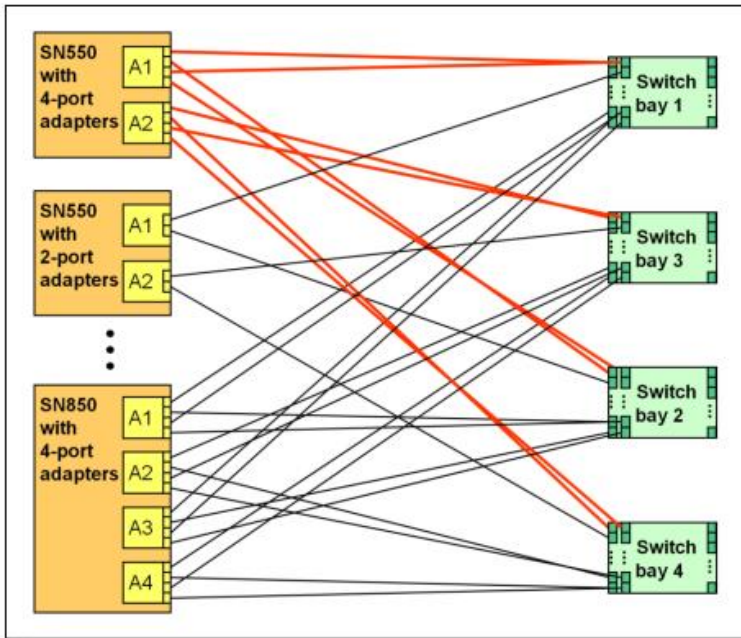


Figure 9. Logical layout of the interconnects between I/O adapters and I/O modules

Source: <https://lenovopress.lenovo.com/lp1043.pdf>

[Lenovo XClarity Administrator](#) >
[Managing servers](#)

Language:

Using remote control

From the Lenovo XClarity Administrator web interface, you can open a remote-control session to a managed server as if you were at a local console. You can use the remote-control session to perform operations such as powering on or off the server, and logically mounting a local or remote drive.

Source:

https://sysmgmt.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.lxca.doc%2Fremotecontrol_thinksystem_use.html/

Note: The XClarity Administrator does not modify the security settings or cryptographic settings (cryptographic mode and the mode used for secure communications) during the management process. You can modify the cryptographic settings after the system is managed (see [Configuring cryptography settings](#)).

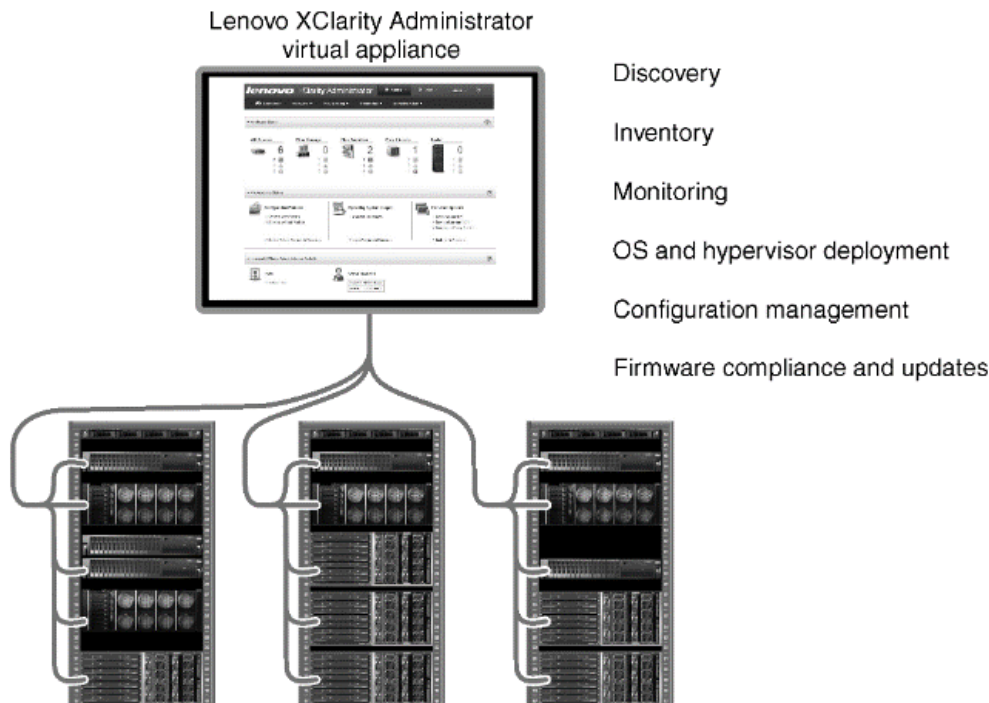
Source:

https://sysmgmt.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.lxca.doc%2Fsetup_manage_systems.html

SSL is a security protocol that provides communication privacy. SSL enables client/server applications to communicate in a way that prevents eavesdropping, tampering, and message forgery. You can configure the XClarity Controller to use SSL support for different types of connections, such as secure web server (HTTPS), secure LDAP connection (LDAPS), CIM over HTTPS, and SSH server, and to manage the certificates that are required for SSL.

Source:

https://sysmgmt.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.systems.management.xcc.doc%2FN1ia_c_configuringsecurity.html

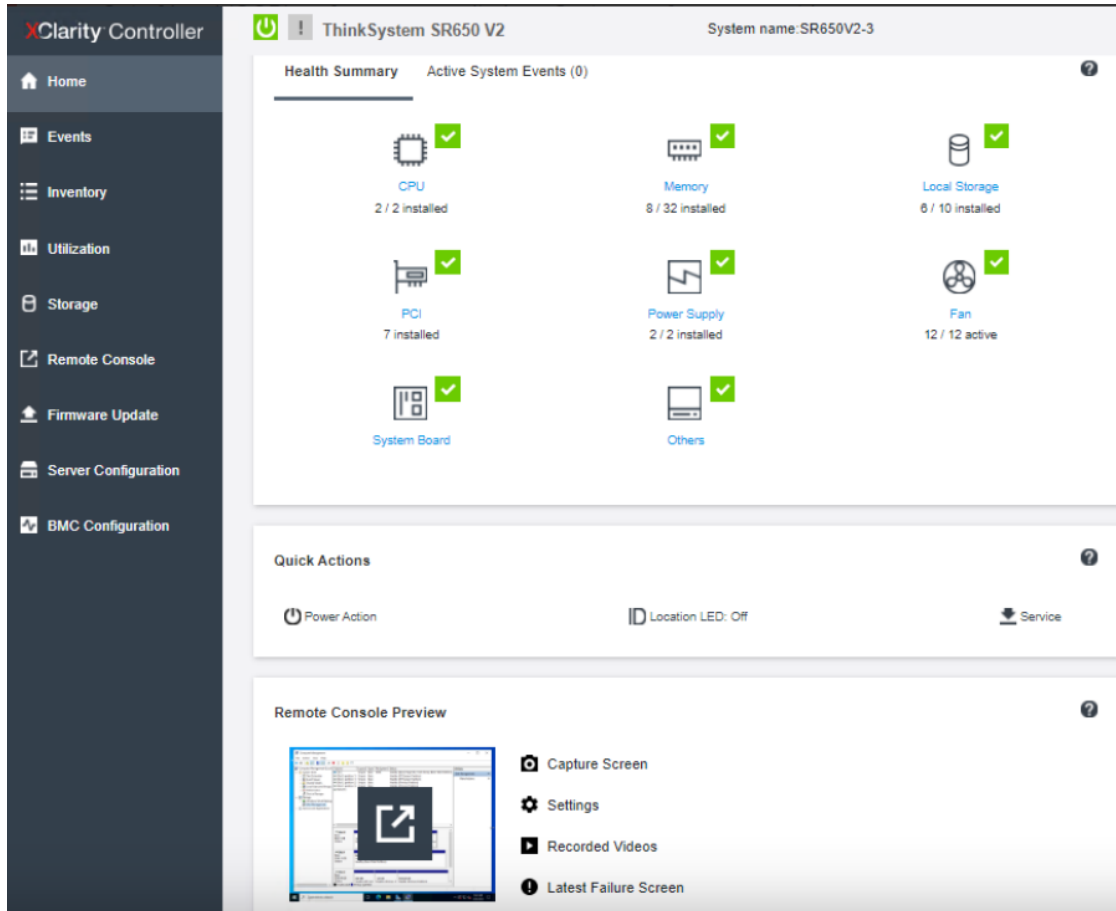


Source:

https://sysmgmt.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.systems.management.xcc.doc%2FN1ia_c_configuringethernet.html

Grant users role-based access and authenticate user credentials.
Establish highly secure communications with managed endpoints using
NIST SP 800-131A and FIPS 140-2 cryptographic standards.

Source: <https://www.lenovo.com/in/en/data-center/software/systems-management/XClarity-Administrator/p/WMD00000366?orgRef=https%253A%252F%252Fwww.google.com%252F>



Source: <https://lenovopress.lenovo.com/lp0880-xcc-support-on-thinksystem-servers>

87. The Accused '016 Products further comprise a first bus. For instance, the ThinkSystem devices include a PCH controller connected to the XCC via a bus, as illustrated below:

System architecture

The following figure shows the architectural block diagram of the SN550, showing the major components and their connections.

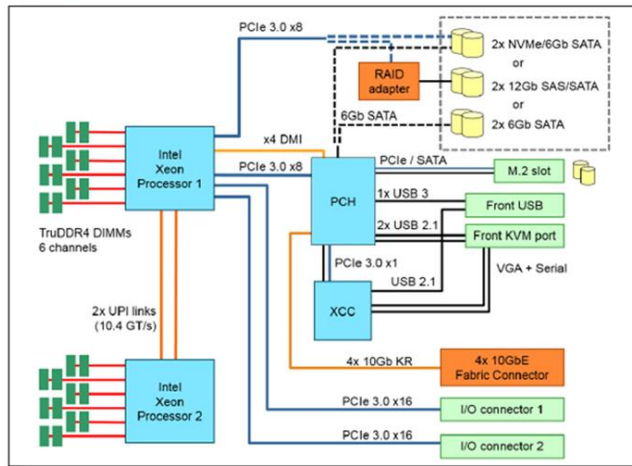


Figure 4. SN550 system architectural block diagram

Source: <https://lenovopress.lenovo.com/lp1043.pdf>

The embedded 10Gb controller is based on the Intel Ethernet Connection X722 network controller which is part of the Intel C624 "Lewisburg" PCH chipset of the SN550 and other Lenovo ThinkSystem servers.

Source: <https://lenovopress.lenovo.com/lp1043.pdf>

The SN550 includes an embedded 4-port 10Gb Intel controller built into the system board. As listed in the [Models](#) section, some SN550 models include the Fabric Connector needed to connect the embedded controller to the midplane of the Flex System chassis. For models that do not include the Fabric Connector, it can be ordered and

Source: <https://lenovopress.lenovo.com/lp1043.pdf>

88. The Accused '016 Products further comprise a bus controller coupled to the processor via the first bus. For example, as discussed above and reiterated below, the ThinkSystem devices include a controller (PCH) connected via a bus to the processor (XCC):

The embedded 10Gb controller is based on the Intel Ethernet Connection X722 network controller which is part of the Intel C624 "Lewisburg" PCH chipset of the SN550 and other Lenovo ThinkSystem servers.

Table 1-2. Intel® C620 Series Chipset SKU Differentiation

SKU	C621	C622	C624	C625	C626	C627	C628	C629
Feature								
Legacy 10/100/1000 Mbps Ethernet	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
LAN Port 0	1G	10/1G	10/1G	10/1G	10/1G	10/1G	10/1G	10/1G
LAN Port 1	1G	10/1G	10/1G	10/1G	10/1G	10/1G	10/1G	10/1G
LAN Port 2	1G	1G	10/1G	10/1G	10/1G	10/1G	10/1G	10/1G
LAN Port 3	1G	1G	10/1G	10/1G	10/1G	10/1G	10/1G	10/1G
Dedicated PCIe Uplink	x1	x8	x16	x16	x16	X16	X16	X16
Muxed x8 PCIe Uplink	Disabled	Disabled	Disabled	Disabled	Enabled	Enabled	Enabled	Enabled
Intel QAT Engines	No	No	No	1	2	3	3	3
Intel QAT Clock Speed	N/A	N/A	N/A	533 MHz	533 MHz	685 MHz	685 MHz	800 MHz
AVID Enabled	No	No	No	Yes (0.85-1.0V)	Yes (0.85-1.0V)	Yes (0.85-1.0V)	Yes (0.85-1.0V)	Yes (0.85-1.0V)

The PCH provides a System Management Bus 2.0 host controller as well as an SMBus Slave Interface. The PCH is also capable of operating in a mode in which it can communicate with I²C compatible devices.

The host SMBus controller supports up to 100 kHz clock speed.

Source: <https://www.intel.com/content/dam/www/public/us/en/documents/datasheets/c620-series-chipset-datasheet.pdf>

The SMBus host controller is used to send commands to other SMBus slave devices. Software sets up the host controller with an address, command, and, for writes, data and optional PEC; and then tells the controller to start. When the controller has finished transmitting data on writes, or receiving data on reads, it generates an SMI# or interrupt, if enabled.

The host SMBus controller supports up to 100 kHz clock speed and is clocked by the RTC clock.

The PCH can perform SMBus messages with either Packet Error Checking (PEC) enabled or disabled. The actual PEC calculation and checking is performed in SW. The SMBus host controller logic can automatically append the CRC byte if configured to do so.

The SMBus host controller is used to send commands to other SMBus slave devices. Software sets up the host controller with an address, command, and, for writes, data and optional PEC; and then tells the controller to start. When the controller has finished transmitting data on writes, or receiving data on reads, it generates an SMI# or interrupt, if enabled.

Source: <https://www.intel.com/content/dam/www/public/us/en/documents/datasheets/c620-series-chipset-datasheet.pdf>

89. In the Accused '016 Products, the bus controller is also coupled to a second bus of the network device that is distinct from the first bus. For instance, the ThinkSystem device's PCH

is connected to the 4x 10GbE fabric connector (or equivalent network interface) via a second bus distinct from the bus connecting the PCH with the XCC, as can be seen below:

System architecture

The following figure shows the architectural block diagram of the SN550, showing the major components and their connections.

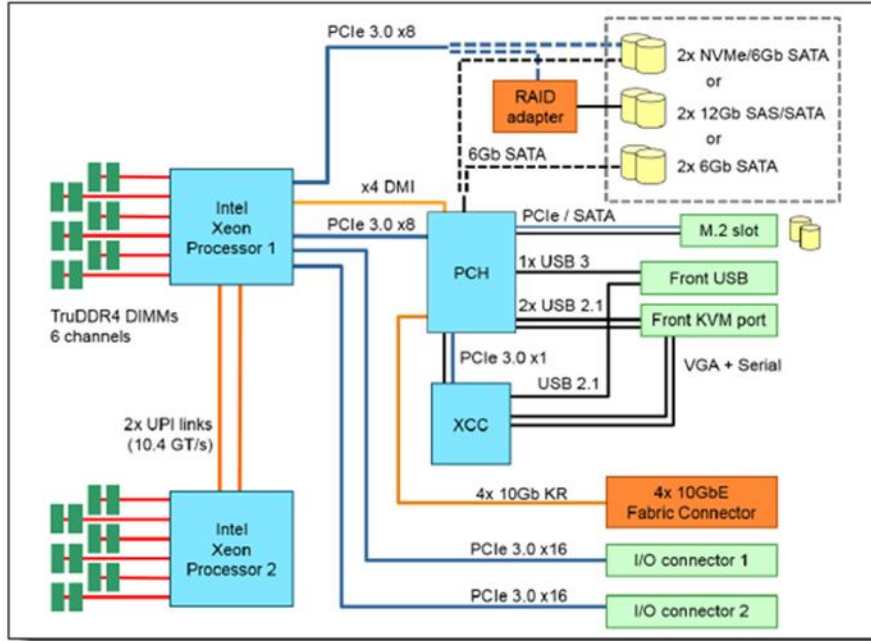


Figure 4. SN550 system architectural block diagram

Source: <https://lenovopress.lenovo.com/lp1043.pdf>

The embedded 10Gb controller is based on the Intel Ethernet Connection X722 network controller which is part of the Intel C624 "Lewisburg" PCH chipset of the SN550 and other Lenovo ThinkSystem servers.

Source: <https://lenovopress.lenovo.com/lp1043.pdf>

The SN550 includes an embedded 4-port 10Gb Intel controller built into the system board. As listed in the Models section, some SN550 models include the Fabric Connector needed to connect the embedded controller to the midplane of the Flex System chassis. For models that do not include the Fabric Connector, it can be ordered and

Source: <https://lenovopress.lenovo.com/lp1043.pdf>

90. In the Accused '016 Products, the bus controller is configured to receive the encrypted form of the management requests from the second bus, and to convey the encrypted form of the management requests to the processor via the first bus. For example, the ThinkSystem device's PCH receives management requests from a remote administrator—directed to the XCC

processor—via the 4x 10GbE fabric connector (or equivalent network interface), which are encrypted for security purposes, as illustrated below:

System architecture

The following figure shows the architectural block diagram of the SN550, showing the major components and their connections.

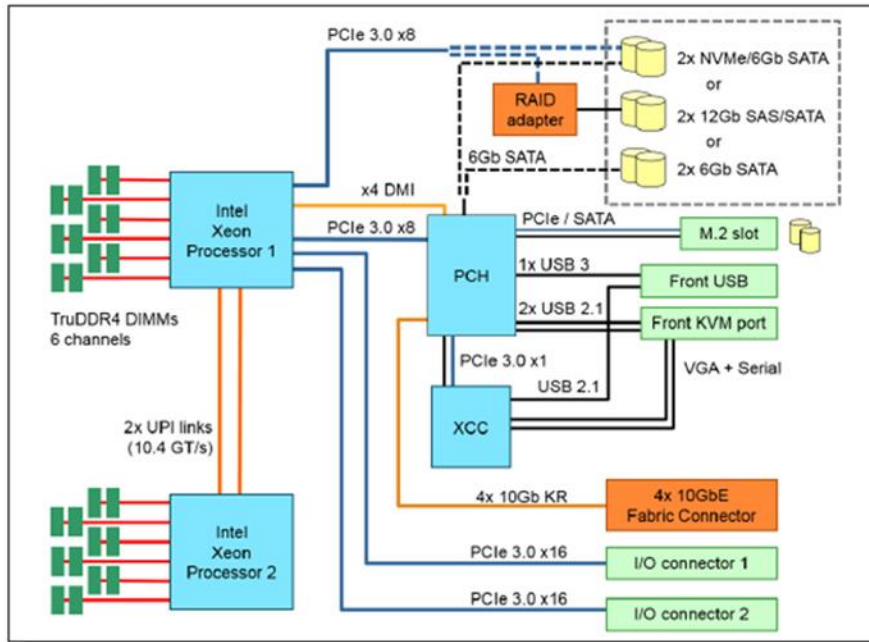


Figure 4. SN550 system architectural block diagram

Source: <https://lenovopress.lenovo.com/lp1043.pdf>

SSL is a security protocol that provides communication privacy. SSL enables client/server applications to communicate in a way that prevents eavesdropping, tampering, and message forgery. You can configure the XClarity Controller to use SSL support for different types of connections, such as secure web server (HTTPS), secure LDAP connection (LDAPS), CIM over HTTPS, and SSH server, and to manage the certificates that are required for SSL.

Source:

https://sysmgt.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.systems.management.xcc.doc%2FN1ia_c_configuringsecurity.html

Grant users role-based access and authenticate user credentials.
Establish highly secure communications with managed endpoints using
NIST SP 800-131A and FIPS 140-2 cryptographic standards.

Source: <https://www.lenovo.com/in/en/data-center/software/systems-management/XClarity-Administrator/p/WMD00000366>

Part number	Description	E		1S Intel				4S Intel				Dense/ Blade				
		SE350 (7Z46/7D1X)	ST50 (7Y48/7Y50)	ST250 (7Y45/7Y46)	SR150 (7Y54)	SR250 (7Y51/7Y52)	SR850 (7X18/7X19)	SR850P (7D2F/2D2G)	SR860 (7X69/7X70)	SR950 (7X11/12/13)	SR850 V2 (7D31/32/33)	SR860 V2 (7Z59/7Z60)	SD530 (7X21)	SD650 (7X58)	SN550 (7X16)	SN650 (7X15)
Integrated LOM for blade servers																
None	Integrated 4-port 10 Gb (requires Fabric Connector)	N	N	N	N	N	N	N	N	N	N	N	N	N	Y	Y

Source: <https://lenovopress.lenovo.com/lp0654-intel-x722-integrated-controller>

Lenovo XClarity Controller (XCC) is an all-new embedded management engine common in every ThinkSystem server.

Virtual presence (remote control) and virtual media capability also come standard in the SN550. The remote control functions include the following:

- Remotely viewing video with graphics resolutions up to 1600x1200 at 75 Hz with up to 32 bits per pixel, regardless of the system state
- Remotely accessing the server using the keyboard and mouse from a remote client
- Capturing blue-screen errors
- International keyboard mapping support
- LDAP-based authentication

Source: <https://www.bechtle.com/shop/medias/5c6c21474c2f85188bee81d2.pdf>

Configuring cryptography settings

Cryptographic management is composed of communication modes and protocols that control the way that secure communications are handled between Lenovo XClarity Administrator and the managed systems.

About this task

The *cryptographic mode* determines how secure communications are handled between XClarity Administrator and all managed systems. If secure communications are implemented, it sets the encryption-key lengths to be used.

Note: Regardless of the cryptography mode that you select, NIST-approved Digital Random Bit Generators are always used, and only 128-bit or longer keys are used for symmetric encryption.

Source:

https://sysmgt.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.lxca.doc%2FCryptography_setmode.html

Access the web page in the LXCA WebUI by clicking **Administration** → **Security** → **Cryptography** as shown in Figure 6.

Figure 6 Administration > Security > Cryptography

Source: <https://lenovopress.lenovo.com/lp1260.pdf>

After the servers are managed by XClarity Administrator, Lenovo XClarity Administrator polls each managed server periodically to collect information, such as inventory, vital product data, and status. You can view and monitor each managed server and perform management actions (such as configuring system settings, deploying operating-system images, and powering on and off).

Note: The XClarity Administrator does not modify the security settings or cryptographic settings (cryptographic mode and the mode used for secure communications) during the management process. You can modify the cryptographic settings after the server is managed (see [Configuring cryptography settings](#)).

Source:

https://sysmgmt.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.lxca.doc%2Fserver_manage.html

91. In the Accused '016 Products, the processor is configured to decrypt the encrypted form of the management requests, and the network device includes a processor configured to facilitate operation of the network device wherein the processor of the apparatus is distinct from the processor included in the network device. For instance, the XCC is a secure management chip embedded in the ThinkSystem device which services management requests from remote

administrators, among other things, that are encrypted. The requests are decrypted at the XCC processor, which is distinct from the device's main processors (in the example below the two Intel Xeon processors):

System architecture

The following figure shows the architectural block diagram of the SN550, showing the major components and their connections.

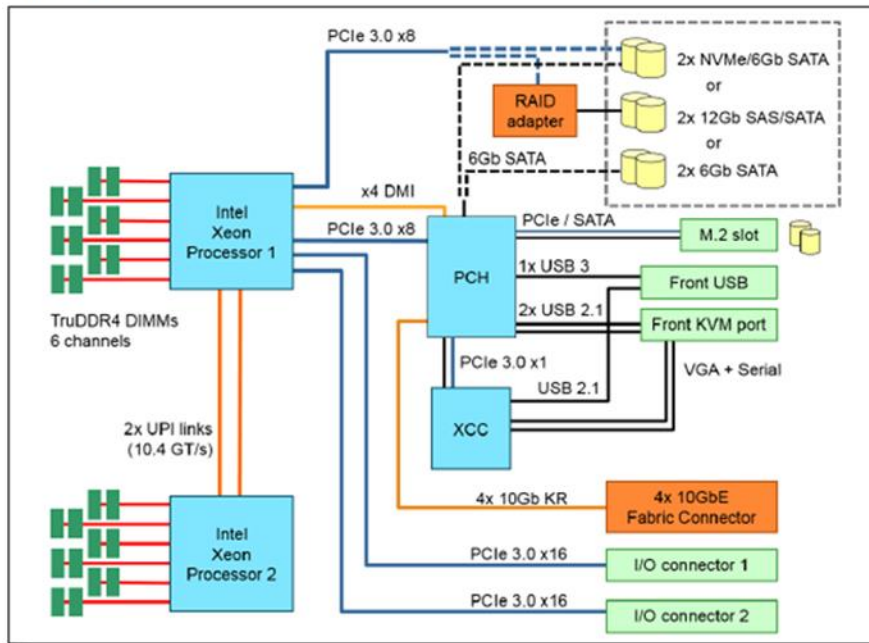


Figure 4. SN550 system architectural block diagram

Source: <https://lenovopress.lenovo.com/lp1043.pdf>

SSL is a security protocol that provides communication privacy. SSL enables client/server applications to communicate in a way that prevents eavesdropping, tampering, and message forgery. You can configure the XClarity Controller to use SSL support for different types of connections, such as secure web server (HTTPS), secure LDAP connection (LDAPS), CIM over HTTPS, and SSH server, and to manage the certificates that are required for SSL.

Source:

https://sysmgmt.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.systems.management.xcc.doc%2FNFI1ia_c_configuringsecurity.html

Grant users role-based access and authenticate user credentials.
Establish highly secure communications with managed endpoints using
NIST SP 800-131A and FIPS 140-2 cryptographic standards.

Source: <https://www.lenovo.com/in/en/data-center/software/systems-management/XClarity-Administrator/p/WMD00000366>

You can use SSL with a self-signed certificate or with a certificate that is signed by a third-party certificate authority. Using a self-signed certificate is the simplest method for using SSL; but, it does create a small security risk. The risk arises because the SSL client has no way of validating the identity of the SSL server for the first connection that is attempted between the client and server. For example, it is possible that a third party might impersonate the XClarity Controller web server and intercept data that is flowing between the actual XClarity Controller web server and the user's web browser. If, at the time of the initial connection between the browser and the XClarity Controller, the self-signed certificate is imported into the certificate store of the browser, all future communications will be secure for that browser (assuming that the initial connection was not compromised by an attack).

Source:

https://sysmgt.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.systems.management.xcc.doc%2FNFIa_c_ssloverview.html

The XClarity Controller contains a hardware random number generator which provides entropy for generating strong cryptographic keys. It has been found that if random numbers are requested too quickly in succession or too soon after device startup, the numbers returned are predictable. To have sufficient entropy:

- There must be no less than two microseconds between requests
- At least the first 32 values generated after startup must be discarded

The current XCC firmware does not include these protections. This may reduce the strength and security of keys generated by the XCC, such as TLS self-signed certificate keys (used to protect https-based communications such as web UI, Redfish, and CIM), ephemeral web UI session keys, ephemeral https session keys, SSH host keys, and so on.

Updating the firmware is strongly recommended. All keys created post-update will be generated with appropriately random data and thus will be secure.

Source: https://support.lenovo.com/in/en/product_security/ps500148-lenovo-xclarity-controller-xcc-risk-of-low-entropy

Cryptographic management is composed of communication modes and protocols that control the way that secure communications are handled between XClarity Administrator and managed systems. The cryptographic mode sets the encryption-key lengths that are used if secure communications are implemented.

XClarity Administrator supports the following SSL/TLS protocols for secure communications:

▶ Legacy

This option is the default option and it enables older cryptographic protocols. When this option is selected, TLS 1.0, TLS 1.1, and TLS 1.2 are supported. If you select this option, you must select the Compatibility cryptographic mode.

▶ TLS 1.2 Server and Client

This option enforces TLS 1.2 cryptography protocols on XClarity Administrator and all managed endpoints. It is set automatically if you choose NIST SP 800-131A for the cryptographic mode.

Source: <https://lenovopress.lenovo.com/sg248296.pdf>

Configure LDAP

If your organization uses an LDAP server, such as Microsoft Active Directory, for authentication and authorization, you can configure XCC to use your LDAP server to authenticate and authorize XCC users too.

Be sure to select Enable Secure LDAP as well. Note that in order to enable it, a valid SSL certificate must first be in place and at least one SSL client trusted certificate must be imported. The LDAP server must also support TLS 1.2 because it is used by LDAP client in XCC.

Source: <https://lenovopress.lenovo.com/lp1260.pdf>

On Security, is where we configure security properties for the XClarity Controller. Here we can also find the Quick Link menu, and we have the option to move over SSL, SSH, IPMI, SYS FW, TPM/TCM, SKLM, and SPM.

The screenshot shows the XClarity Controller web interface for a ThinkSystem SR950 server. The left sidebar contains navigation options: Home, Events, Inventory, Utilization, Remote Console, Firmware Update, Server Configuration, BMC Configuration, Backup and Restore, License, Network, Security, and User/LDAP. The main content area displays several security settings:

- SSL Certificate Management:** A signed certificate is installed. Expiration: April 29, 2029 6:54 PM.
- SSH Server:** A SSH server key is installed. Status: On.
- IPMI over KCS Access:** Enabled.
- Prevent System Firmware Down-Level:** Disabled.
- Assert Physical Presence:** De-assert.

Source: <https://www.storagereview.com/review/lenovo-xclarity-controller-xcc-review>

Processor	One or two second-generation Intel Xeon Processor Scalable Family of processors (formerly codename "Cascade Lake"). Supports processors with up to 28 cores, core speeds up to 3.8 GHz, and TDP ratings up to 165W.
-----------	---

Source: <https://lenovopress.lenovo.com/lp1043.pdf>

Configuring cryptography settings

Cryptographic management is composed of communication modes and protocols that control the way that secure communications are handled between Lenovo XClarity Administrator and the managed systems.

About this task

The *cryptographic mode* determines how secure communications are handled between XClarity Administrator and all managed systems. If secure communications are implemented, it sets the encryption-key lengths to be used.

Note: Regardless of the cryptography mode that you select, NIST-approved Digital Random Bit Generators are always used, and only 128-bit or longer keys are used for symmetric encryption.

Source:

https://sysmgmt.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.lxca.doc%2FCryptography_setmode.html

92. Additionally, Defendant has been, and currently is, an active inducer of infringement of the '016 patent under 35 U.S.C. § 271(b) and a contributory infringer of the '016 patent under 35 U.S.C. § 271(c).

93. Defendant has actively induced, and continues to actively induce, infringement of the '016 patent by causing others to use, offer for sale, or sell in the United States, products or services covered by the '016 patent, including the Accused '016 Products. Defendant provides these products and services to others, such as customers, resellers, partners, and end-users, who, in turn, use, provision for use, offer for sale, or sell those products and services, which directly infringe the '016 patent. Defendant's inducement includes the directions and instructions found at one or more of the following links, the provision of which was on-going as of the filing of the Complaint in case 6:23-cv-0068-ADA, and remains on-going as of the filing of this Complaint, and the content of which is specifically illustrated above:

- <https://lenovopress.lenovo.com/lp1043.pdf>
- <https://lenovopress.lenovo.com/lp1397-thinksystem-sn550-v2-server>
- <https://lenovopress.lenovo.com/lp1396-thinksystem-sd650-n-v2-server>
- <https://lenovopress.com/lp0637-thinksystem-sn550-server-xeon-sp-gen-1>
- <https://lenovopress.com/lp0880-xcc-support-on-thinksystem-servers>

- https://sysmgt.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.lxca.doc%2Fremot_econtrol_thinksystem_use.html
- https://sysmgt.lenovofiles.com/help/index.jsp?topic=%2Fllxcc_frontend%2Fllxcc_overview.html&cp=3
- https://sysmgt.lenovofiles.com/help/topic/com.lenovo.systems.management.xcc.doc/NN1ia_c_configuringethernet.html?cp=3_0_3_1_0
- https://sysmgt.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.systems.management.xcc.doc%2Fdw1lm_c_accessingtheimmwebinterface.html&cp=3_0_2_0
- https://sysmgt.lenovofiles.com/help/topic/com.lenovo.systems.management.xcc.doc/dw1lm_t_settinguptheimmnetworkconnection.html?cp=3_0_2_0_0
- https://sysmgt.lenovofiles.com/help/topic/com.lenovo.systems.management.xcc.doc/NN1ia_c_configuringethernet.html
- https://sysmgt.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.lxca.doc%2Fsecurity_implement.html
- https://sysmgt.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.lxca.doc%2Faccess_control_setspecificdevices.html
- https://sysmgt.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.systems.management.xcc.doc%2FNN1ia_c_configuringDNS.html
- https://sysmgt.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.lxca.doc%2Fsecurity_implement.html
- https://sysmgt.lenovofiles.com/help/topic/com.lenovo.lxca.doc/lxca_qsg_install_kvm.pdf
- <https://www.lenovo.com/in/en/data-center/software/systems-management/c/systems-management>
- <https://www.lenovo.com/us/en/data-center/software/management/>
- https://systemx.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.systemx.common.nav.doc%2Foverview_imm2.html&cp=0_4
- https://systemx.lenovofiles.com/help/topic/com.lenovo.sysx.imm2.doc/dw1lm_c_ch3_configuringtheimm.html
- https://systemx.lenovofiles.com/help/topic/com.lenovo.sysx.imm2.doc/NN1ia_c_configuringnetworkprotocolproperties.html

94. Defendant has contributed to, and continues to contribute to, the infringement of the '016 patent by others by knowingly providing one or more components, for example the XCC processor and PCH included in the Accused Products, a portion thereof, and/or the software/hardware modules responsible for the accused functionality described herein, that, when

installed, configured, and used result in systems that, as intended by Defendant described above, directly infringe one or more claims of the '016 patent.

95. Defendant knew of the '016 patent, or should have known of the '016 patent, but was willfully blind to its existence. Upon information and belief, Defendant had actual knowledge of the '016 patent at least as early as February 2, 2023, the date the Complaint in case 6:23-cv-0068-ADA was filed or as early as February 3, 2023, the date upon which IV notified Defendant of the filing of the Complaint in the aforementioned action. Alternatively, upon information and belief, Defendant has had actual knowledge of the '016 patent since receipt of this Complaint or service on Defendant of this Complaint.

96. By the time of trial, Defendant will or should have known and intended (since receiving such notice) that its continued actions would infringe and would actively induce and contribute to the infringement of the '016 patent.

97. Defendant has committed, and continues to commit, contributory infringement by selling products and services that directly infringe the '016 patent when used by a third party, such as the Accused '016 Products, and that are a material part of the invention, knowing them to be especially made or adapted for use in infringement of the '016 patent and not staple articles or commodities of commerce suitable for substantial non-infringing use.

98. As a result of Defendant's acts of infringement, IV has suffered and will continue to suffer damages in an amount to be determined at trial.

COUNT III

(Defendant's Infringement of U.S. Patent No. 7,089,443)

99. The preceding paragraphs are reincorporated by reference as if fully set forth herein.

100. The '443 patent claims and teaches, *inter alia*, an improved way to provide for future frequency increases, while maintaining a synchronous design methodology and exploiting the trend towards making functional blocks more autonomous.

101. The inventions improved upon then-existing microprocessor design by creating a MCD microarchitecture which uses a globally asynchronous locally synchronous clocking style. This microarchitecture afforded a number of technical solutions to unsolved, technical problems, with notable advantages over a singly and globally clocked design, more fully discussed below.

102. The inventions claimed in the '443 patent allowed for more autonomous functional blocks operating under more independent local domain clocks, which implies less onerous global clock distribution requirements, permitting potentially higher frequencies within each domain and greater scalability in future process generations.

103. The inventions claimed in the '443 patent further allowed for the designers of each domain to no longer be as constrained by the speeds of critical paths in other domains, affording them greater freedom in each domain to optimize the tradeoffs among clock speed, latency, and the exploitation of application parallelism via complex hardware structures.

104. The inventions claimed in the '443 patent further allowed for the use of separate voltage inputs, external voltage regulators, and controllable clock frequency circuits in each clock domain, which in turned allowed for finer-grained dynamic voltage and frequency scaling, and thus lower energy, than could be achieved with single clock, single-core-voltage systems.

105. The inventions claimed in the '443 patent further allowed for the ability to dynamically alter the clock speed in each domain, which in turn meant that the clock-rate tradeoff could be tailored to application characteristics within each individual domain, thereby improving both performance and energy efficiency.

106. More specifically, the claims of the '443 patent recite a multiple-clock-domain microprocessor. The system covered by the asserted claims includes a plurality of domains. It further includes, for each of the plurality of domains, a clock for separately generating a clock signal at a frequency for that domain, the frequency being dynamically changeable independently of the frequencies of the clock signals generated for others of the plurality of domains. And it includes, for each of the plurality of domains, a voltage input for receiving a voltage which is dynamically changeable independently of the voltages applied to said others of the plurality of domains.

107. The system covered by the asserted claims, therefore, differs markedly from the prior systems in use at the time of this invention, which lacked the claimed combination of the plurality of domains with clocks generating clock signals at frequencies for each of the domains, wherein the frequency is dynamically changeable independently of the frequencies for the clock signals generated in the other domains; and the plurality of domains with voltage input for receiving voltage which is dynamically changeable independently of the voltages applied in the other domains.

108. Defendant has directly infringed, and continues to directly infringe, at least claim 1 of the '443 patent by making, using, testing, selling, offering for sale, and importing products and services covered by the '443 patent. Defendant's products and services that infringe the '443 patent include all products and services that include Arm Cortex-Axx processors, such as the Lenovo Yoga Tablet 3, Lenovo Tab P11, Lenovo Tab P11 Pro, Lenovo Tab P11 Pro Gen 2, Lenovo Tab M10, Lenovo ThinkPad X13s, Motorola Edge Plus, Motorola G Power, and any other Lenovo devices that include an ARM-based processor with substantially the same functionality as described below, (together the "Accused '443 Products" or "Accused Products").

109. Claim 1 of the '443 patent is reproduced below:

1. A multiple clock domain microprocessor comprising:

a plurality of domains;

for each of the plurality of domains, a clock for separately generating a clock signal at a frequency for that domain, the frequency being dynamically changeable independently of the frequencies of the clock signals generated for others of the plurality of domains; and

for each of the plurality of domains, a voltage input for receiving a voltage which is dynamically changeable independently of the voltages applied to said others of the plurality of domains.

110. The Accused '443 Products each provide an MCD microprocessor. As one example, the Accused '443 Products are products that include a MCD microprocessor, such as the Lenovo Tab P11 Pro, which includes an Octa-Core Processor with a combination of Arm Cortex A-76 and Arm Cortex A-55 cores in a scalable DynamIQ big.LITTLE configuration, as seen below:

Home > Tablets > Android > Lenovo Tab Series > Tab P11 Pro

READY TO SHIP

Tab P11 Pro Tablet + pen + keyboard bundle

★★★★★ 4.4 (321) Part Number: ZA7C0124US

Save \$80.00 **13% off**
Est Value ~~\$599.99~~

\$519.99 OR **\$87/mo** suggested payments w/6 mo promo financing
[See How](#) | [Prequalify](#)

Add To Cart

📦 **Delivery FREE** Standard Delivery: Get it by Thu, Jan 26
[Delivery options for 60654](#)

Special Offers

Business Price: 🛒 Members Only [Join LenovoPRO & Save](#)

Student & Teachers Price: 🎓 Verify in Cart to Save [Learn More](#)

🎁 Earn Double Rewards = \$31 + Get it by Thu, Jan 26 [Join Now!](#)

Compare

System Specs: [View All Models >](#)

Processor
Qualcomm® Snapdragon™ 730G Octa-Core Processor (8 x Kryo 470 CPU, up to 2.2 GHz)

Operating System
Android™ 10

Source: <https://www.lenovo.com/us/en/p/tablets/android-tablets/lenovo-tab-series/lenovo-tb-j706/za7c0124us>

Snapdragon 730	
General Info	
Designer	Qualcomm, ARM Holdings
Manufacturer	Samsung
Model Number	SDM730
Market	Mobile
Introduction	April 9, 2019 (announced) April 9, 2019 (launched)
General Specs	
Family	Snapdragon 700
Frequency	2,200 MHz, 1,800 MHz
Microarchitecture	
ISA	ARMv8 (ARM)
Microarchitecture	Cortex-A76, Cortex-A55
Core Name	Kryo 470 Gold, Kryo 470 Silver
Process	8 nm
Technology	CMOS
Word Size	64 bit
Cores	8
Threads	8
Max Memory	8 GiB
Multiprocessing	
Max SMP	1-Way (Uniprocessor)
Succession	
← Snapdragon 710	

Source:

https://en.wikichip.org/wiki/qualcomm/snapdragon_700/730#:~:text=Snapdragon%20730%20is%20a%20mid,Gold%20operating%20at%202.2%20GHz

Specifications

The Arm Cortex-A76 CPU delivers laptop-class performance with smartphone efficiency, bringing the same experience to all classes of intelligent mobile compute devices.

The second generation premium core built on DynamIQ technology. Paired with a Cortex-A55 CPU in a scalable DynamIQ big.LITTLE configuration, Cortex-A76 delivers laptop-class performance with mobile efficiency, bringing the mobile experience (fast responsiveness, always on, always connected) into all classes of intelligent mobile compute devices. With superior energy efficiency and far greater single-threaded performance, Cortex-A76 CPU extends battery life and improves user experience for sustained high performance across even the most complex compute tasks.

Source: <https://developer.arm.com/ip-products/processors/cortex-a/cortex-a76>

A1.1 About the core

The Cortex-A76 core is a high-performance and low-power Arm product that implements the Armv8-A architecture.

The Cortex-A76 core supports:

- The Armv8.2-A extension.
- The RAS extension.
- The Load acquire (LDAPR) instructions introduced in the Armv8.3-A extension
- The Dot Product support instructions introduced in the Armv8.4-A extension.
- The PSTATE *Speculative Store Bypass Safe* (SSBS) bit and the speculation barriers (CSDB, SSBB, PSSBB) instructions introduced in the Armv8.5-A extension.

The Cortex-A76 core has a *Level 1* (L1) memory system and a private, integrated *Level 2* (L2) cache. It also includes a superscalar, variable-length, out-of-order pipeline.

The Cortex-A76 core is implemented inside the *DynamIQ Shared Unit* (DSU) cluster. For more information, see the *Arm® DynamIQ™ Shared Unit Technical Reference Manual*.

Source: <https://documentation-service.arm.com/static/602fa9141e2cbd4091013c48?token=>

A1.1 About the core

The Cortex-A55 core is a mid-range, low-power core that implements the ARMv8-A architecture with support for the v8.2 extension, the RAS extension, the Load acquire (LDAPR) instructions introduced in the ARMv8.3 extension, and the Dot Product instructions introduced in the ARMv8.4 extension.

The core has a *Level 1* (L1) memory system, and private *Level 2* (L2) cache. The core is implemented inside the DynamIQ Shared Unit (DSU) as a Little core and is highly configurable with other cores.

Source: <https://documentation-service.arm.com/static/5e7e1405b471823cb9de57ae?token=>


Feature	Cortex-A55	Cortex-A76
Architecture	Armv8.2-A	Armv8.2-A (AArch32 at EL0 only)

Source: <https://developer.arm.com/documentation/102826/latest/>

111. Furthermore, the Accused '443 Products each comprise a MCD microprocessor with plurality of domains. For example, the Accused Products include a MCD microprocessor with multiple domains, such as any one of the individual cores and its CPU bridge (e.g., Core 0 and its CPU bridge) from the DynamIQ cluster (“User Core”) in combination with the DynamIQ cluster shared logic/unit (“DSU”) including the L3 cache and snoop filters, which collectively comprise an MCD microprocessor. That MCD processor has a domain that includes Core 0 and another domain that includes the DSU.

The *DynamIQ™ Shared Unit-110* (DSU-110) provides a shared L3 memory system, snoop control and filtering, and other control logic to support a cluster of A-class architecture cores. The cluster is called the DSU-110 DynamIQ™ cluster. Additionally, all the external interfaces to *System on Chip* (SoC) are provided through the DSU-110.

Figure 2-1: DSU-110 DynamIQ™ cluster



A DSU-110 DynamIQ™ cluster consists of between one and eight cores, with up to three different types of cores in the same cluster. Cores can be configured for various performance points during macrocell implementation and run at different frequencies and voltages.

Source: <https://documentation-service.arm.com/static/611e9446d5c3af0155491bf8?token=>



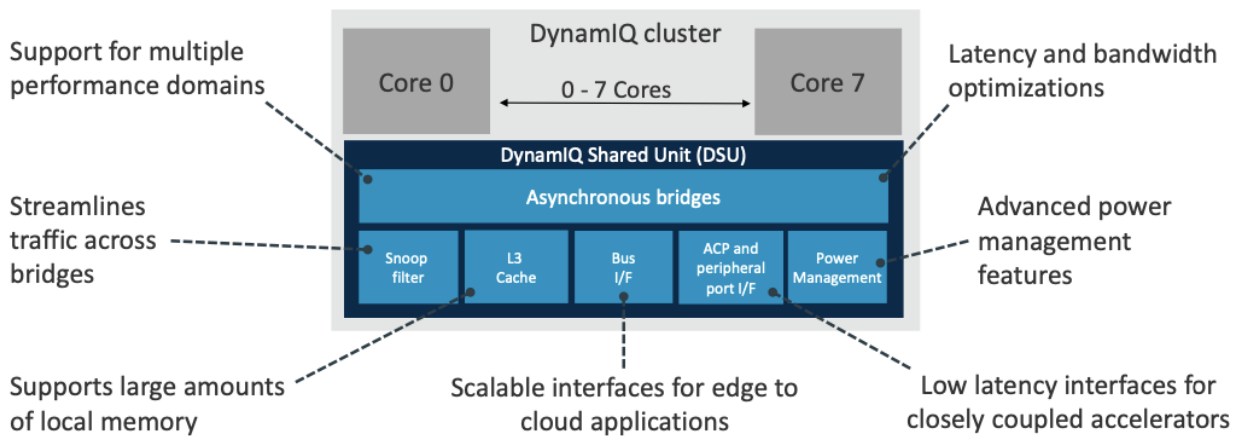


Figure 3-2: DynamIQ™ cluster shared logic components



Source: <https://documentation-service.arm.com/static/611e9446d5c3af0155491bf8?token=>

DynamIQ Shared Unit (DSU)




Source: https://old.hotchips.org/wp-content/uploads/hc_archives/hc29/HC29.22-Tuesday-Pub/HC29.22.80-Architectdure-Pub/HC29,22,820-DynamIQ-Greenhaigh-ARM.pdf

112. The Accused '443 Products further include, for each of the plurality of domains, a clock for separately generating a clock signal at a frequency for that domain. For example, there is a clock for separately generating a clock signal at a frequency for any one of the User Cores (e.g., Core 0 and its CPU bridge) of the DynamIQ cluster, and another clock for separately generating another clock signal at another frequency for the DSU.

The *DynamiQ™ Shared Unit-110* (DSU-110) has multiple clock domains. Each core or complex can be implemented in a separate clock domain.

The following figure shows the clock domains for an example cluster with four standalone cores.

Figure 4-1: DSU-110 clock domains



The cluster contains several clock domains for functionality that is likely to be connected to different clocks in the system. Within each core, the CPU bridge contains asynchronous bridges for all crossings between the core and cluster clock domains. Each CPU bridge is split, with one half of each bridge in the core clock domain and the other half in the relevant cluster domain. At the cluster level, there is the *Snoop Control Unit* (SCU) bridge which contains crossings between the cluster clock domains as required.

4.1 Clocks

The *DynamiQ™ Shared Unit-110* (DSU-110) has a separate clock signal for each standalone core or complex. There are also separate clocks for the internal logic, and some of the external interfaces.

The following table describes the clock signals of the DSU-110.

Table 4-1: DSU-110 clock signals

Signal	Description
CORExCLK	The clocks for each of the cores in the cluster that are not part of a complex. x is the core instance number, for example, CORE0CLK is the clock for core 0. These signals clock all core logic, including L1 and L2 caches.
COMPLEXxCLK	The clocks for each complex in the cluster. Each clock is connected to all cores in the respective complex. x is the complex instance number, for example, COMPLEX0CLK is the clock for complex 0.
SCLK	This clock is used for the <i>Snoop Control Unit</i> (SCU), L3 memory system, and all the external interfaces, including AXI, CHI, and <i>Accelerator Coherency Port</i> (ACP). It is also used for cores that are configured to run synchronously with the DSU-110.

Source: <https://documentation-service.arm.com/static/611e9446d5c3af0155491bf8?token=>

113. In the Accused '443 Products, the frequency of the clock signal of a domain is dynamically changeable independent of the frequencies of the clock signals generated for other domains. For example, the frequency of the clock signal of any one of the User Cores of the DynamIQ cluster is dynamically changeable independent of the frequencies of the clock signals of the DSU.

Clock management

Clock gating is supported through Q-Channel requests from an external clock controller to the DSU-110. The Q-Channels allow individual control of the following clock input signals:

- **ATCLK**
- **CORExCLK** where x is the core instance number
- **COMPLEXxCLK** where x is the complex instance number
- **GICCLK**
- **PCLK**
- **PERIPHCLK**
- **PPUCLK**
- **SCLK**

All clocks can be driven fully asynchronously to each other. The DSU-110 contains all the necessary synchronizing logic for crossing between clock domains. There are no clock dividers and no latches in the design. The entire design is rising-edge triggered.

While there is no functional requirement for any of the clocks to have any relationship to any of the others, the DSU-110 is designed with the following expectations to achieve an acceptable performance:

- The **CORExCLK** or **COMPLEXxCLK** can be dynamically scaled to match the performance requirements of that core.
- **SCLK** is recommended to run between the maximum **CORExCLK** or **COMPLEXxCLK** frequency and approximately half of the maximum **CORExCLK** or **COMPLEXxCLK** frequency.

L3 cache system can be clocked at a rate synchronous to the external system interconnect or at integer multiples

Source: <https://documentation-service.arm.com/static/611e9446d5c3af0155491bf8?token=>

3.1.1 Integration of the cores in the cluster

When you implement a DSU-110 DynamIQ™ cluster, all interfacing between the cores, complexes, and the *DynamIQ™ Shared Unit-110* (DSU-110) is implemented automatically. All the external signal inputs and outputs pass through the DSU-110. The DSU-110 buffers and resynchronizes many of these signals to allow cores to be clocked at different speeds.

The memory interfacing of each core is internally connected to the DSU-110 L3 memory system. Where necessary, the DSU-110 implements additional buffering to compensate for different clock rates of the core and DSU-110 L3 memory system.

Each core has an external clock interface, which is routed through the DSU-110 to the respective core.

Power management and clock gating

The DebugBlock implements two Q-Channel interfaces, one for requests to gate the PCLK clock, and a second for requests to control the Debug power domain.

4.1 Clocks

The *DynamIQ™ Shared Unit-110* (DSU-110) has a separate clock signal for each standalone core or complex. There are also separate clocks for the internal logic, and some of the external interfaces.

Cluster features

The DSU-110 has the following cluster features:

- Support for cores running independently at different frequencies and voltages known as *Dynamic Voltage Frequency Scaling* (DVFS). For cores in a complex, DVFS is only possible for the whole complex not for individual cores.

Source: <https://documentation-service.arm.com/static/611e9446d5c3af0155491bf8?token=>

A4.8 Clock, voltage, and power domains

The DynamIQ cluster microarchitecture supports multiple clock, voltage, and power domains.

The number of domains that are implemented depends on the choices made by the SoC implementer. There might be fewer in your SoC.

The following diagram shows the clock, voltage, and power domains supported by the DSU and cores.

- Voltage domains are indicated by dashed outlines.
- Blocks that are in the same power domain have the same color.

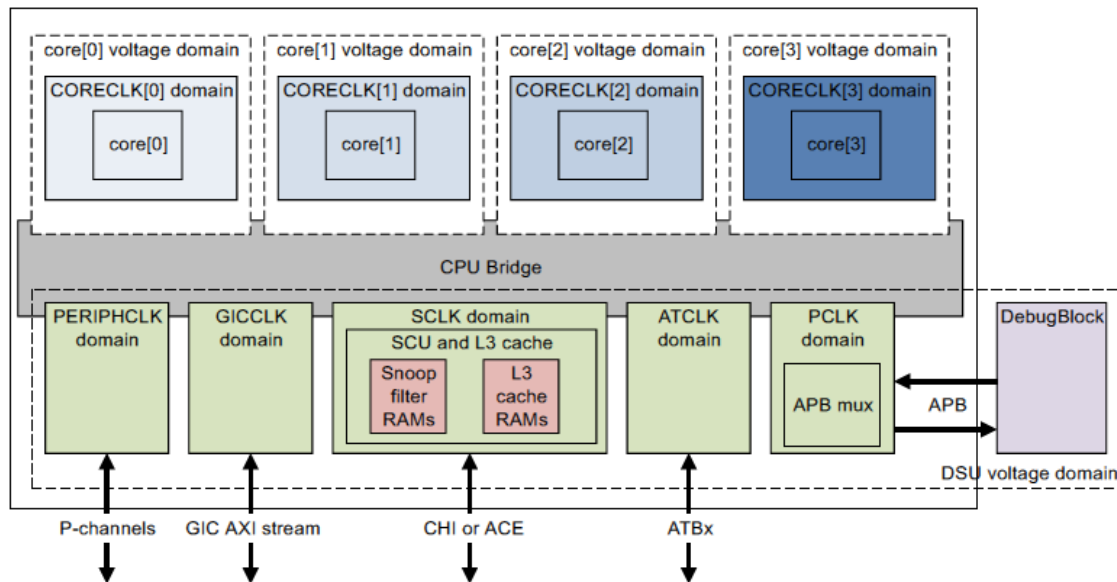


Figure A4-2 DSU Clock, voltage, and power domains

Clock domains

Each core can be implemented in a separate clock domain. The DSU has multiple clock domains.

The CPU Bridge contains all asynchronous bridges for crossing clock domains, and is split with one half of each bridge in the core clock domain and the other half in the relevant cluster domain. Each core can be implemented with or without an asynchronous bridge. If the asynchronous bridge is not implemented, the core is in the SCLK clock domain.

Source: <https://documentation-service.arm.com/static/611e9446d5c3af0155491bf8?token=>

The following diagram shows a cluster that is composed of two sets of cores in a DynamIQ big.LITTLE configuration.




Figure A1-1 DynamIQ cluster

Within the DSU, are the L3 cache, the *Snoop Control Unit* (SCU), internal interfaces to the cores, and external interfaces to the SoC.

- The shared L3 cache simplifies process migration between the cores.

————— **Note** —————

Some cores can be configured without L2 caches. To these cores, the shared L3 cache appears as an L2 cache. The term 'L3 cache' is used throughout this document to describe the shared cache.

- The *Snoop Control Unit* (SCU) maintains coherency between caches in the cores and L3. The SCU includes a Snoop Filter to optimize coherency maintenance operations.
- Internal interfaces to the cores are configured during macrocell implementation and are not directly visible.
- External interfaces are connected to the SoC.

Each core can be configured either to be run synchronously with the DSU, sharing the clock, or asynchronously, with an independent clock.

Source: <https://documentation-service.arm.com/static/5e7e16e0b2608e4d7f0a3030>

A3.1 Clocks

The DSU requires clock signals for each of the cores, internal logic, and external interfaces.

The following table describes the clocks.

Table A3-1 DSU clock signals

Signal	Description
CORECLK[CN:0]	The per-core clocks for all core logic including L1 and L2 caches.
SCLK	The clock for the SCU and L3 memory system, including the ACE or CHI master interface. SCLK is also used for any cores that are configured to run synchronously to the DSU.
PCLK	The clock for the DebugBlock and DSU debug APB interfaces. <p style="text-align: center;">————— Note —————</p> The DebugBlock and cluster both have PCLK inputs. You might choose to connect these to the same clock. Alternatively, you might choose to place an asynchronous bridge between the two clock inputs, in which case they might be different clocks.
ATCLK	The clock for the ATB trace buses output from the DSU. <p style="text-align: center;">————— Note —————</p> All ATB buses output from the DSU share the same clock.
GICCLK	The clock for the GIC AXI-stream interface between the DSU and an external GIC.
PERIPHCLK	The clock for peripheral logic inside the DSU such as timers, and clock and power management logic.

All clocks can be driven fully asynchronously to each other. The DSU contains all the necessary synchronizing logic for crossing between clock domains. There are no clock dividers and no latches in the design. The entire design is rising edge triggered.

Source: <https://documentation-service.arm.com/static/5e7e16e0b2608e4d7f0a3030>

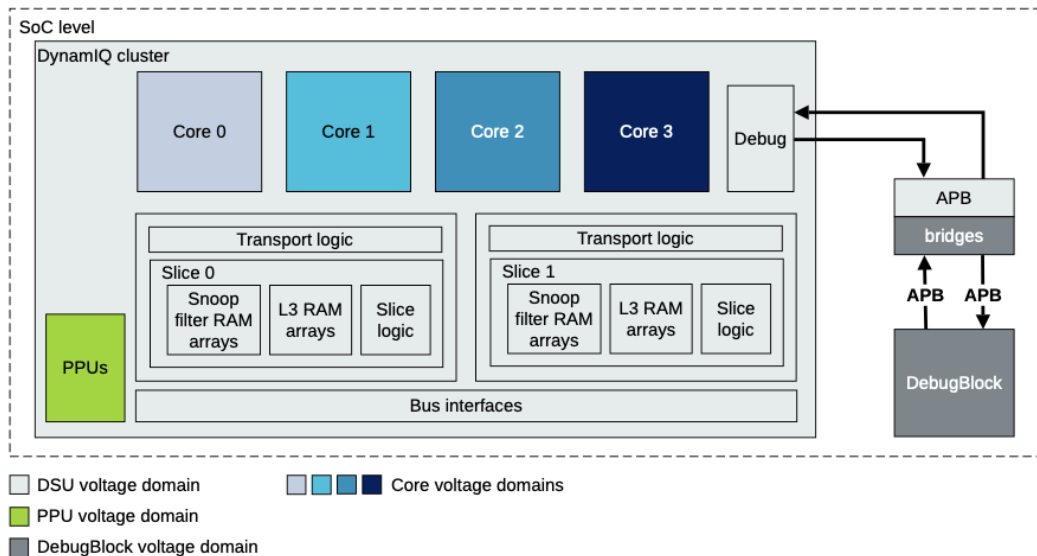
ARM DynamIQ

- ARM DynamIQ technology
 - Different microarchitectures can exist within same cluster.
 - Share the same last level of cache (L3), i.e. DSU (DynamIQ shared unit).
 - Different compute capacity and frequency domains for big and LITTLE.
- All the entities can do DVFS: big, LITTLE and DSU.
- Same voltage domain possible for DSU and big or/and LITTLE CPUs.
- DSU controls the cache bandwidth available to CPUs.
- DSU bandwidth configured based on requirements from CPUs.

Source: http://retis.sssup.it/luca/ospm-summit/2018/Downloads/Device_notifications_and_improved_efficiency.pdf

114. The Accused '443 Products further include, for each of the plurality of domains, a voltage input for receiving a voltage. For example, any one of the User Cores of the DynamIQ cluster and the DSU, each have a voltage input for receiving a voltage.

Figure 5-6: DSU-110 voltage domains



Having each core in a separate voltage domain allows *Dynamic Voltage Frequency Scaling* (DVFS) to be applied to each core.

Cluster features

The DSU-110 has the following cluster features:

- Support for Arm[®]v9.0-A architecture cores
- Support for up to three types of core, and a maximum of eight cores in the cluster
- *Power Policy Units* (PPUs) providing autonomous power management of the L3 cache and the cores
- Support for cores running independently at different frequencies and voltages known as *Dynamic Voltage Frequency Scaling* (DVFS). For cores in a complex, DVFS is only possible for the whole complex not for individual cores.

Source: <https://documentation-service.arm.com/static/611e9446d5c3af0155491bf8?token=>

115. In the Accused '443 Products, for each of the voltage inputs for each of the plurality of domains, the voltage is dynamically changeable independent of the voltages applied to said others of the plurality of domains. For example, any one of the User Cores of the DynamIQ cluster and the DSU, each have their own independent voltage domains which are dynamically changeable independent of the voltages applied to the other voltage domains.

A4.1 About DSU power management

The DSU supports a range of low-power modes and cache RAM powerdown modes.

The DSU supports the following power modes:

On

On mode is the normal mode of operation where all the core and DSU functionality is available. The DSU individually disables internal clocks, and inputs to unused functional blocks. Only the logic that is in use consumes dynamic power.

Functional retention

Functional retention allows the L3 cache and snoop filter RAMs to be put temporarily in to a retention state while the L3 cache is not being accessed. The contents of the cache RAMs are retained.

Memory retention

Memory retention mode allows the L3 cache and snoop filter RAMs to be held in retention while the rest of the cluster is powered down. Keeping the RAMs in retention reduces the energy cost of writing dirty lines back to memory and reduces the cluster response time on powerup. It is not possible to snoop the cache in this mode, so it is important that no other external coherent agents are active (for example, cores external to the cluster, or other coherent devices). In practice, this mode can only be used in a coherent system when the cluster is the only active agent.

Off

In off mode, power is removed completely, and no state is retained. To avoid losing data, the cores within the cluster, and the cluster itself must first be taken out of coherence.

The DSU supports clock, voltage, and power domains that can be controlled by external logic. The cluster, in conjunction with power management software, gives operating requirement hints to an external power controller. The power controller is responsible for coordinating power management with the rest of the SoC, switching and isolating power and voltage domains, and controlling clock gating cells.

Source: <https://documentation-service.arm.com/static/5e7e16e0b2608e4d7f0a3030>

A4.8 Clock, voltage, and power domains

The DynamIQ cluster microarchitecture supports multiple clock, voltage, and power domains.

The number of domains that are implemented depends on the choices made by the SoC implementer. There might be fewer in your SoC.

The following diagram shows the clock, voltage, and power domains supported by the DSU and cores.

- Voltage domains are indicated by dashed outlines.
- Blocks that are in the same power domain have the same color.

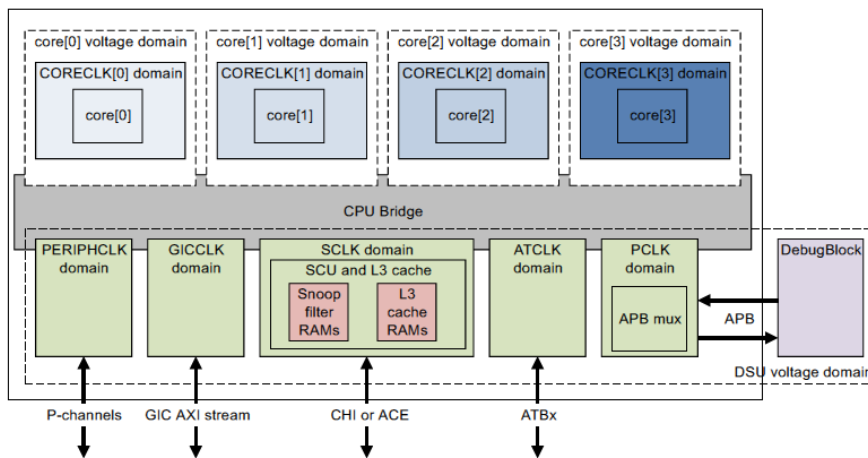


Figure A4-2 DSU Clock, voltage, and power domains

Voltage domains


Each core can be implemented in a separate voltage domain. The DSU has a single separate voltage domain, allowing, for example, the DSU to be in the same voltage domain as the SoC interconnect and other system components.

Source: <https://documentation-service.arm.com/static/5e7e16e0b2608e4d7f0a3030>

ARM DynamIQ

- ARM DynamIQ technology
 - Different microarchitectures can exist within same cluster.
 - Share the same last level of cache (L3), i.e. DSU (DynamIQ shared unit).
 - Different compute capacity and frequency domains for big and LITTLE.
- All the entities can do DVFS: big, LITTLE and DSU.
- Same voltage domain possible for DSU and big or/and LITTLE CPUs.
- DSU controls the cache bandwidth available to CPUs.
- DSU bandwidth configured based on requirements from CPUs.

Source: http://retis.sssup.it/luca/ospm-summit/2018/Downloads/Device_notifications_and_improved_efficiency.pdf



Source: <https://vdocuments.net/enabling-arm-dynamiq-dynamiq-introduction-dynamiq-and-arm-trusted-firmware.html?page=14>

116. Additionally, Defendant has been, and currently is, an active inducer of infringement of the '443 patent under 35 U.S.C. § 271(b) and a contributory infringer of the '443 patent under 35 U.S.C. § 271(c).

117. Defendant has actively induced, and continues to actively induce, infringement of the '443 patent by causing others to use, offer for sale, or sell in the United States, products or services covered by the '443 patent, including but not limited to the '443 Accused Products and any other products or services that include ARM-based processors with the functionality described above. Defendant provides these products and services to others, such as customers, resellers, partners, and end-users, who, in turn, use, provision for use, offer for sale, or sell those products and services, which directly infringe the '443 patent as described above. Defendant's inducement includes the directions and instructions found at one or more of the following links, the provision of which has been on-going as of the filing of the First Amended Complaint in case 6:23-cv-0068-ADA and much of the content of which is specifically illustrated above:

- <https://www.lenovo.com/us/en/p/tablets/android-tablets/lenovo-tab-series/lenovo-tab-p11-plus/wmd00000476>
- https://psref.lenovo.com/syspool/Sys/PDF/Lenovo_Tablets/Tab_P11_Plus/Tab_P11_Plus_Spec.pdf
- <https://vaosia.com/en-de/products/lenovo-pro-11-5-touchscreen-2-in-1-notebook-arm-cortex-a76-octa-core-2-20-ghz>
- <https://www.lenovo.com/us/en/p/tablets/android-tablets/lenovo-tab-series/tab-p11-pro-gen-2/len10310011?IPromoId=LEN999994>
- https://psref.lenovo.com/syspool/Sys/PDF/Lenovo_Tablets/Tab_P11_Pro_2nd_Gen/Tab_P11_Pro_2nd_Gen_Spec.pdf
- <https://www.lenovo.com/us/en/p/tablets/android-tablets/lenovo-tab-series/lenovo-tab-m10-plus-gen-3/len10310010>
- https://psref.lenovo.com/syspool/Sys/PDF/Lenovo_Tablets/Tab_M10_Plus_3rd_Gen/Tab_M10_Plus_3rd_Gen_Spec.pdf
- <https://news.lenovo.com/pressroom/press-releases/first-thinkpad-powered-by-snapdragon-multi-day-battery-life-ai-5g/>

- <https://developer.arm.com/Processors/Cortex-A76>
- <https://www.arm.com/products/silicon-ip-cpu/cortex-a/cortex-a76>
- <https://documentation-service.arm.com/static/611e9446d5c3af0155491bf8?token=>
- <https://www.youtube.com/watch?v=oHyfOv8DN0w>
- <https://www.youtube.com/watch?v=uULKRRQnDPI>
- <https://www.youtube.com/watch?v=UklxqU2jaTg>
- <https://www.youtube.com/watch?v=c7Rz5s3IPK4>
- <https://www.youtube.com/watch?v=ScTu-x2vI4>

118. Defendant has contributed to, and continues to contribute to, the infringement of the '443 patent by others by knowingly providing one or more components, for example the ARM-Cortex-Axx-based CPU included in the Accused Products, a portion thereof, and/or the software/hardware modules responsible for the accused functionality described herein, that, when installed, configured, and used result in systems that, as intended by Lenovo described above, directly infringe one or more claims of the '443 patent.

119. Defendant knew of the '443 patent, or should have known of the '443 patent, but was willfully blind to its existence. Upon information and belief, Defendant had actual knowledge of the '443 patent since at least as early as the filing of the First Amended Complaint in case 6:23-cv-0068-ADA, or alternatively, at least as early as Defendant's receipt of this Complaint. Alternatively, upon information and belief, Defendant has had knowledge of the '443 patent since the service upon Defendant of the Complaint in this action.

120. By the time of trial, Defendant will or should have known and intended (since receiving such notice) that its continued actions would infringe and would actively induce and contribute to the infringement of the '443 patent.

121. Defendant has committed, and continues to commit, contributory infringement by selling products and services that directly infringe the '443 patent when used by a third party, such

as the Accused '443 Products, and that are a material part of the invention, knowing them to be especially made or adapted for use in infringement of the '443 patent and not staple articles or commodities of commerce suitable for substantial non-infringing use.

122. As a result of Defendant's acts of infringement, IV has suffered and will continue to suffer damages in an amount to be determined at trial.

COUNT IV

(Defendant's Infringement of U.S. Patent No. 7,623,439)

123. The preceding paragraphs are reincorporated by reference as if fully set forth herein.

124. The '439 patent claims and teaches, *inter alia*, an improved signal transmitting system capable of manipulating OFDM data packets and data streams using improved cyclic diversity schemes, thereby improving packet reception performance when compared to conventional packet diversity mechanisms by reducing packet error rates, among other benefits.

125. The inventions improved upon then-existing cyclic diversity schemes in wireless communication by enabling a cyclic diversity scheme by which a portion of an OFDM packet's symbol data is cyclically advanced into the guard interval with respect to the original OFDM signal and then each signal is sent to a receiver device at substantially the same time from two respective antennas. This allowed for improved acquisition and correlation at the receiver while at the same time keeping intersymbol interference and unintentional beamforming to a minimum.

126. Unlike in prior art systems and methods, the cyclic diversity taught by the '439 patent uses cyclic advancement as opposed to delay. Doing so substantially reduces the probability of unintentional beamforming.

127. More specifically, one exemplary embodiment comprises an improved cyclic diversity system in which a logic circuit is configured to cyclically advance samples of a symbol

data portion of an OFDM packet to be transmitted on a first antenna, relative to the samples of a symbol data portion of another OFDM packet to be transmitted on another antenna. The duration of the cyclic advance is less than the duration of a guard interval portion of the OFDM packet. By using a cyclic advance as described above, the symbol data portions of the two different transmitted signals are better decorrelated, thus reducing the probability of unintentional beamforming. The performance of wireless networks is thereby improved by the technologies disclosed and claimed in the '439 patent.

128. The system and methods covered by the asserted claims, therefore, differs markedly from prior art systems in use at the time of this invention, which lacked the claimed combination of cyclically advancing a first OFDM packet by shifting the samples in a first direction an amount less than a sample duration of the guard interval portion to generate a shifted version of the first OFDM packet in which at least a non-zero number of samples from the symbol data portion of the first OFDM packet are shifted into the guard interval portion of the shifted version, and a same non-zero number of samples from the guard interval portion of the first OFDM packet are shifted out of the guard interval portion of the shifted version. And, further where both versions are substantially simultaneously transmitted.

129. Defendant has directly infringed and continues to directly infringe at least claim 1 of the '439 patent by making, using, selling, offering for sale, or importing products and/or services covered by the '439 patent. Defendant's products and/or services that infringe the '439 patent include all wireless communication products that support IEEE 802.11n, 802.11ac and 802.11ax, including the transmission of multiple spatial streams, which requires a cyclic diversity shift when transmitting OFDM packets, that are made, used, sold, or offered for sale by or on behalf of

Defendant and/or its subsidiaries or parent companies (cumulatively, “the ’439 Accused Products”), including but not limited to, the Motorola Edge+.

130. Claim 1 of the ’439 patent is reproduced below:

1. A method for transmitting orthogonal frequency division multiplexing (OFDM) signals comprising:

generating a first OFDM packet for transmission including a guard interval portion and a symbol data portion each comprised of a plurality of samples;

cyclically advancing the first OFDM packet by shifting the samples in a first direction an amount less than a sample duration of the guard interval portion to generate a shifted version of the first OFDM packet for transmission in which at least a non-zero number of the samples from the symbol data portion of the first OFDM packet are shifted into the guard interval portion of the shifted version and a same non-zero number of samples from the guard interval portion of the first OFDM packet are shifted out of the guard interval portion of the shifted version; and

substantially simultaneously transmitting the first OFDM packet and the shifted version of the OFDM packet.

131. The ’439 Accused Products are configured to perform a method for transmitting OFDM signals. As one example, the Motorola Edge+ supports the IEEE 802.11n standard, including the transmission of multiple spatial streams, which requires a forward shift diversity feature for transmitting OFDM signals:



5G

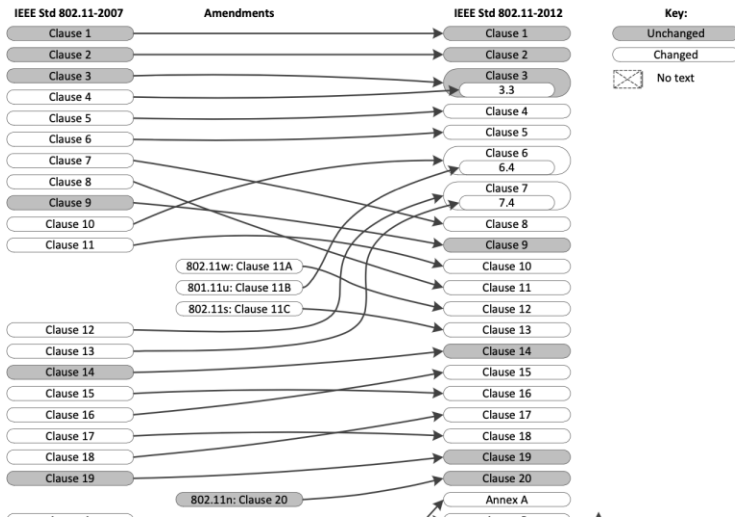
Lightning-fast 5G on the **motorola edge+** is capable of over 4 Gbps[†]— speeds never seen before on a smartphone. Whether browsing, streaming, or gaming, you demand performance. No lag time, no interruptions.

processor & RAM

The Qualcomm® Snapdragon™ 865 Mobile Platform is the world's fastest[†], with an AI engine that can process 15 trillion operations per second. Add to that 12GB of DDR5 memory and 256GB of UFS 3.0 storage[§] for faster bandwidth and lower battery drain. It's all the power you can handle.

- Qualcomm® FastConnect™ 6800 Subsystem
- Wi-Fi Standards: Wi-Fi 6 (802.11ax), 802.11ac Wave 2, 802.11a/b/g/n
- Wi-Fi Spectral Bands: 2.4 GHz, 5 GHz
- Peak speed: 1.774 Gbps
- Channel Utilization: 20/40/80 MHz
- 8-stream sounding (for 8x8 MU-MIMO)
- MIMO Configuration: 2x2 (2-stream)
- MU-MIMO (Uplink & Downlink)
- 1024 QAM (2.4 & 5 GHz)
- OFDMA (2.4 and 5 GHz)
- Dual-band simultaneous (DBS)
- Wi-Fi Security: WPA3-Enterprise, WPA3-Enhanced Open, WPA3 Easy Connect, WPA3-Personal

Source: <https://www.motorola.com/us/smartphones-motorola-edge-plus/p>



Source: IEEE 802.11-2012.

20. High Throughput (HT) PHY specification

20.1 Introduction

20.1.1 Introduction to the HT PHY

Clause 20 specifies the PHY entity for a high throughput (HT) orthogonal frequency division multiplexing (OFDM) system.

In addition to the requirements found in Clause 20, an HT STA shall be capable of transmitting and receiving frames that are compliant with the mandatory PHY specifications defined as follows:

- In Clause 18 when the HT STA is operating in a 20 MHz channel width in the 5 GHz band
- In Clause 17 and Clause 19 when the HT STA is operating in a 20 MHz channel width in the 2.4 GHz band

The HT PHY is based on the OFDM PHY defined in Clause 18, with extensibility up to four spatial streams, operating in 20 MHz bandwidth. Additionally, transmission using one to four spatial streams is defined for operation in 40 MHz bandwidth. These features are capable of supporting data rates up to 600 Mb/s (four spatial streams, 40 MHz bandwidth).

The HT PHY data subcarriers are modulated using binary phase shift keying (BPSK), quadrature phase shift keying (QPSK), 16-quadrature amplitude modulation (16-QAM), or 64-QAM. Forward error correction (FEC) coding (convolutional coding) is used with a coding rate of 1/2, 2/3, 3/4, or 5/6. LDPC codes are added as an optional feature.

Source: IEEE Std 802.11-2012 pp. 1669.

20.1.2 Scope

The services provided to the MAC by the HT PHY consist of two protocol functions, defined as follows:

- a) A PHY convergence function, which adapts the capabilities of the physical medium dependent (PMD) system to the PHY service. This function is supported by the physical layer convergence procedure (PLCP), which defines a method of mapping the PSDUs into a framing format (PPDU) suitable for sending and receiving PSDUs between two or more STAs using the associated PMD system.

Source: IEEE Std 802.11-2012 pp. 1669.

20.1.4 PPDU formats

The structure of the PPDU transmitted by an HT STA is determined by the TXVECTOR FORMAT, CH_BANDWIDTH, CH_OFFSET, and MCS parameters as defined in Table 20-1. The effect of the CH_BANDWIDTH, CH_OFFSET, and MCS parameters on PPDU format is described in 20.2.3.

The FORMAT parameter determines the overall structure of the PPDU as follows:

- *Non-HT format (NON_HT)*: Packets of this format are structured according to the Clause 18 (OFDM) or Clause 19 (ERP) specification. Support for non-HT format is mandatory.
- *HT-mixed format (HT_MF)*: Packets of this format contain a preamble compatible with Clause 18 and Clause 19 receivers. The non-HT-STF (L-STF), the non-HT-LTF (L-LTF), and the non-HT SIGNAL field (L-SIG) are defined so they can be decoded by non-HT Clause 18 and Clause 19 STAs. The rest of the packet cannot be decoded by Clause 18 or Clause 19 STAs. Support for HT-mixed format is mandatory.
- *HT-greenfield format (HT_GF)*: HT packets of this format do not contain a non-HT compatible part. Support for HT-greenfield format is optional. An HT STA that does not support the reception of an HT-greenfield format packet shall be able to detect that an HT-greenfield format packet is an HT transmission (as opposed to a non-HT transmission). In this case, the receiver shall decode the HT-SIG and determine whether the HT-SIG cyclic redundancy check (CRC) passes.

Source: IEEE Std. 802.11-2012, pp. 1669-70.




Figure 20-1—PPDU format

Source: IEEE Std. 802.11-2012, pp. 1682.

The HT portion of the HT-mixed format preamble enables estimation of the MIMO channel to support demodulation of the HT data by HT STAs. The HT portion of the HT-mixed format preamble also includes the HT-SIG field, which supports HT operation. The SERVICE field is prepended to the PSDU.

Source: IEEE Std. 802.11-2012, pp. 1682.

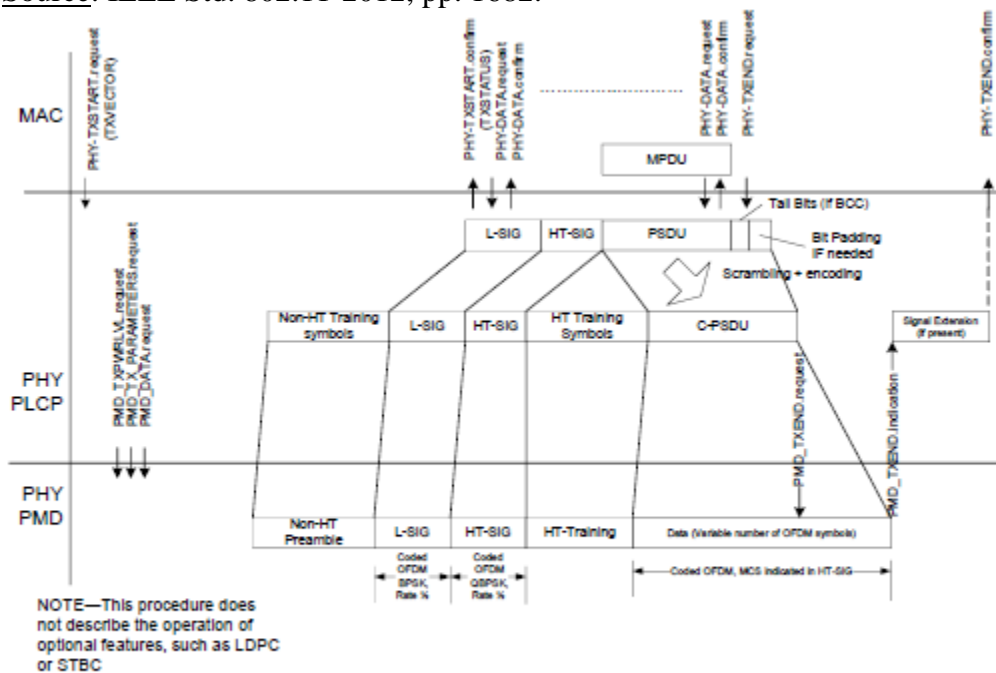


Figure 20-22—PLCP transmit procedure (HT-mixed format PPDU)

Source: IEEE Std. 802.11-2012, p. 1748.

132. The method practiced by the '439 Accused Products includes generating a first OFDM packet for transmission including a guard interval portion and a symbol data portion each comprised of a plurality of samples. For instance, the 802.11 transmitter in the Accused Products creates an OFDM packet, known as an HT-SIG OFDM packet, which includes a symbol data portion comprised of a plurality of samples and a guard interval portion comprised of a plurality of samples as seen below:

20. High Throughput (HT) PHY specification

20.1 Introduction

20.1.1 Introduction to the HT PHY

Clause 20 specifies the PHY entity for a high throughput (HT) orthogonal frequency division multiplexing (OFDM) system.

In addition to the requirements found in Clause 20, an HT STA shall be capable of transmitting and receiving frames that are compliant with the mandatory PHY specifications defined as follows:

- In Clause 18 when the HT STA is operating in a 20 MHz channel width in the 5 GHz band
- In Clause 17 and Clause 19 when the HT STA is operating in a 20 MHz channel width in the 2.4 GHz band

The HT PHY is based on the OFDM PHY defined in Clause 18, with extensibility up to four spatial streams, operating in 20 MHz bandwidth. Additionally, transmission using one to four spatial streams is defined for operation in 40 MHz bandwidth. These features are capable of supporting data rates up to 600 Mb/s (four spatial streams, 40 MHz bandwidth).

The HT PHY data subcarriers are modulated using binary phase shift keying (BPSK), quadrature phase shift keying (QPSK), 16-quadrature amplitude modulation (16-QAM), or 64-QAM. Forward error correction (FEC) coding (convolutional coding) is used with a coding rate of 1/2, 2/3, 3/4, or 5/6. LDPC codes are added as an optional feature.

Source: IEEE Std 802.11-2012 pp. 1669.

20.1.2 Scope

The services provided to the MAC by the HT PHY consist of two protocol functions, defined as follows:

- a) A PHY convergence function, which adapts the capabilities of the physical medium dependent (PMD) system to the PHY service. This function is supported by the physical layer convergence procedure (PLCP), which defines a method of mapping the PSDUs into a framing format (PPDU) suitable for sending and receiving PSDUs between two or more STAs using the associated PMD system.

Source: IEEE Std 802.11-2012 pp. 1669.

20.1.4 PPDU formats

The structure of the PPDU transmitted by an HT STA is determined by the TXVECTOR FORMAT, CH_BANDWIDTH, CH_OFFSET, and MCS parameters as defined in Table 20-1. The effect of the CH_BANDWIDTH, CH_OFFSET, and MCS parameters on PPDU format is described in 20.2.3.

The FORMAT parameter determines the overall structure of the PPDU as follows:

- *Non-HT format (NON_HT)*: Packets of this format are structured according to the Clause 18 (OFDM) or Clause 19 (ERP) specification. Support for non-HT format is mandatory.
- *HT-mixed format (HT_MF)*: Packets of this format contain a preamble compatible with Clause 18 and Clause 19 receivers. The non-HT-STF (L-STF), the non-HT-LTF (L-LTF), and the non-HT SIGNAL field (L-SIG) are defined so they can be decoded by non-HT Clause 18 and Clause 19 STAs. The rest of the packet cannot be decoded by Clause 18 or Clause 19 STAs. Support for HT-mixed format is mandatory.
- *HT-greenfield format (HT_GF)*: HT packets of this format do not contain a non-HT compatible part. Support for HT-greenfield format is optional. An HT STA that does not support the reception of an HT-greenfield format packet shall be able to detect that an HT-greenfield format packet is an HT transmission (as opposed to a non-HT transmission). In this case, the receiver shall decode the HT-SIG and determine whether the HT-SIG cyclic redundancy check (CRC) passes.

Source: IEEE Std. 802.11-2012, pp. 1669-70.

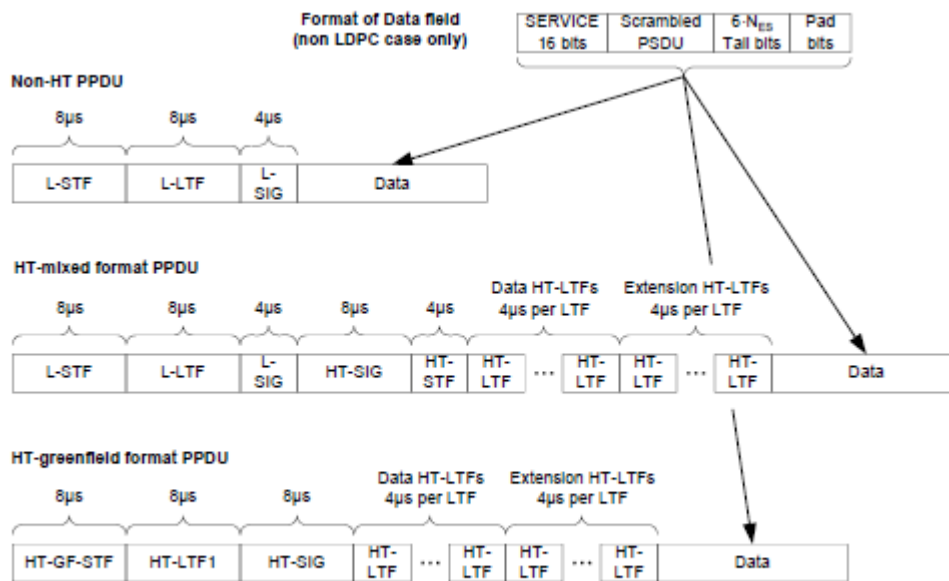


Figure 20-1—PPDU format

Source: IEEE Std. 802.11-2012, pp. 1682.

The HT portion of the HT-mixed format preamble enables estimation of the MIMO channel to support demodulation of the HT data by HT STAs. The HT portion of the HT-mixed format preamble also includes the HT-SIG field, which supports HT operation. The SERVICE field is prepended to the PSDU.

Source: IEEE Std. 802.11-2012, pp. 1682.

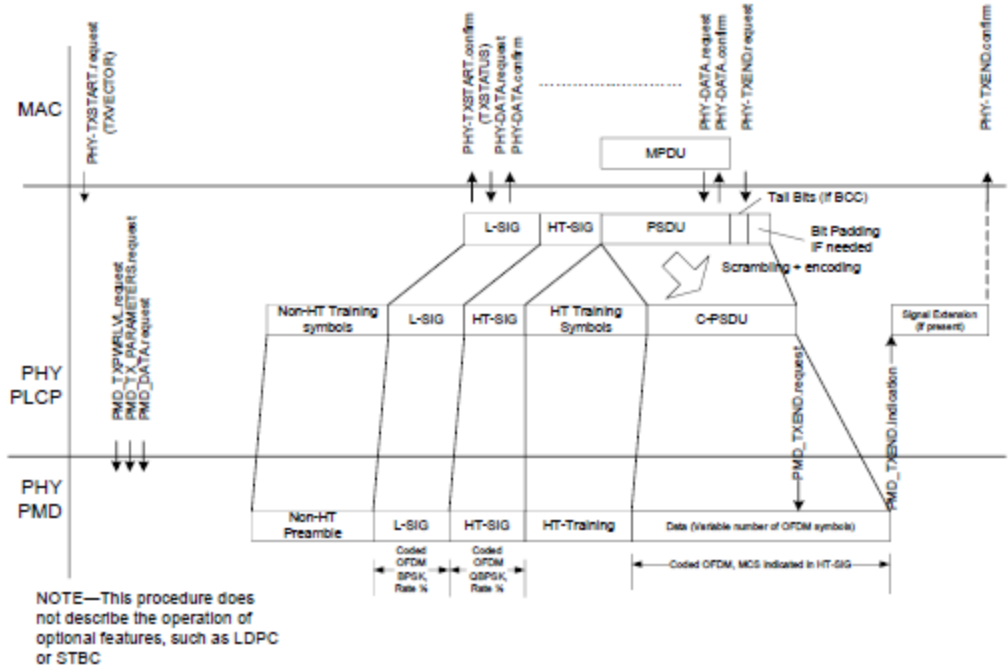


Figure 20-22—PLCP transmit procedure (HT-mixed format PDU)

Source: IEEE Std. 802.11-2012, p. 1748.

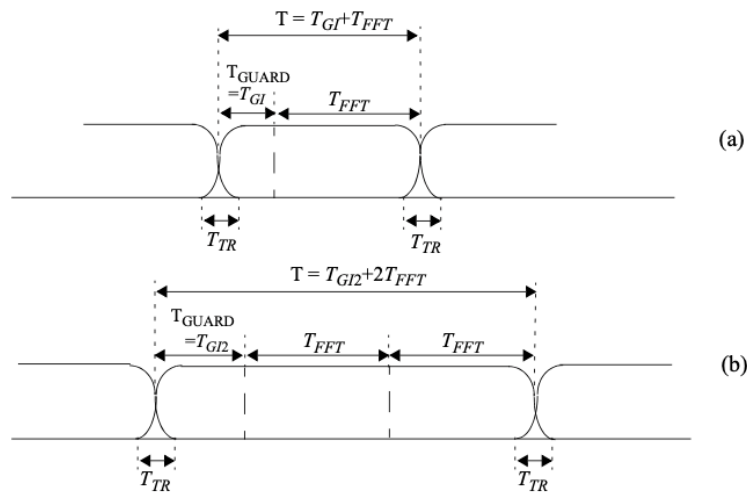


Figure 18-2—Illustration of OFDM frame with cyclic extension and windowing for (a) single reception or (b) two receptions of the FFT period

Source: IEEE Std. 802.11-2012, p. 1592.

T_{FFT} : Inverse Fast Fourier Transform (IFFT) / Fast Fourier Transform (FFT) period	3.2 μs ($1/\Delta_F$)
T_{SIGNAL} : Duration of the SIGNAL BPSK-OFDM symbol	4.0 μs ($T_{GI} + T_{FFT}$)
T_{GI} : GI duration	0.8 μs ($T_{FFT}/4$)


Source: IEEE Std. 802.11-2012, p. 1590-91.

20.3.9.4.3 HT-SIG definition

The HT-SIG is used to carry information required to interpret the HT packet formats. The fields of the HT-SIG are described in Table 20-11.

Table 20-11—HT-SIG fields

Field	Number of bits	Explanation and coding
Modulation and Coding Scheme	7	Index into the MCS table. See NOTE 1.
CBW 20/40	1	Set to 0 for 20 MHz or 40 MHz upper/lower. Set to 1 for 40 MHz.
HT Length	16	The number of octets of data in the PSDU in the range of 0 to 65 535.



Source: IEEE Std. 802.11-2012, p. 1682, 1748, 1699.

The HT-SIG is composed of two parts, HT-SIG₁ and HT-SIG₂, each containing 24 bits, as shown in Figure 20-6. All the fields in the HT-SIG are transmitted LSB first, and HT-SIG₁ is transmitted before HT-SIG₂.

The HT-SIG parts shall be encoded at $R = 1/2$, interleaved, and mapped to a BPSK constellation, and they have pilots inserted following the steps described in 18.3.5.6, 18.3.5.7, 18.3.5.8, and 18.3.5.9, respectively. The BPSK constellation is rotated by 90° relative to the L-SIG in order to accommodate detection of the start of the HT-SIG. The stream of 96 complex numbers generated by these steps is divided into two groups of 48 complex numbers: $d_{k,n}$, $0 \leq k \leq 47$, $n = 0, 1$. The time domain waveform for the HT-SIG in an HT-mixed format packet in a 20 MHz transmission shall be as shown in Equation (20-16).

$$r_{HT-SIG}^{i_{TX}}(t) = \frac{1}{\sqrt{N_{TX} \cdot N_{HT-SIG}^{Tone}}} \sum_{n=0}^1 w_{T_{SYM}}(t - nT_{SYM}) \cdot \sum_{k=-26}^{26} (jD_{k,n} + P_{n+1}P_k) \exp(j2\pi k \Delta_F(t - nT_{SYM} - T_{GI} - T_{CS}^{i_{TX}})) \quad (20-16)$$

Source: IEEE Std. 802.11-2012, p. 1700.

18.3.2.6 Discrete time implementation considerations

The following descriptions of the discrete time implementation are informational.

In a typical implementation, the windowing function is represented in discrete time. As an example, when a windowing function with parameters $T = 4.0 \mu\text{s}$ and a $T_{TR} = 100 \text{ ns}$ is applied, and the signal is sampled at 20 Msample/s, it becomes

$$w_T[n] = w_T(nT_S) = \begin{cases} 1 & 1 \leq n \leq 79 \\ 0.5 & 0, 80 \\ 0 & \text{otherwise} \end{cases} \quad (18-5)$$

Source: IEEE Std. 802.11-2012, p. 1593.

Figure 20-2 and Figure 20-3 show example transmitter block diagrams. In particular, Figure 20-2 shows the transmitter blocks used to generate the HT-SIG of the HT-mixed format PPDU. These transmitter blocks are

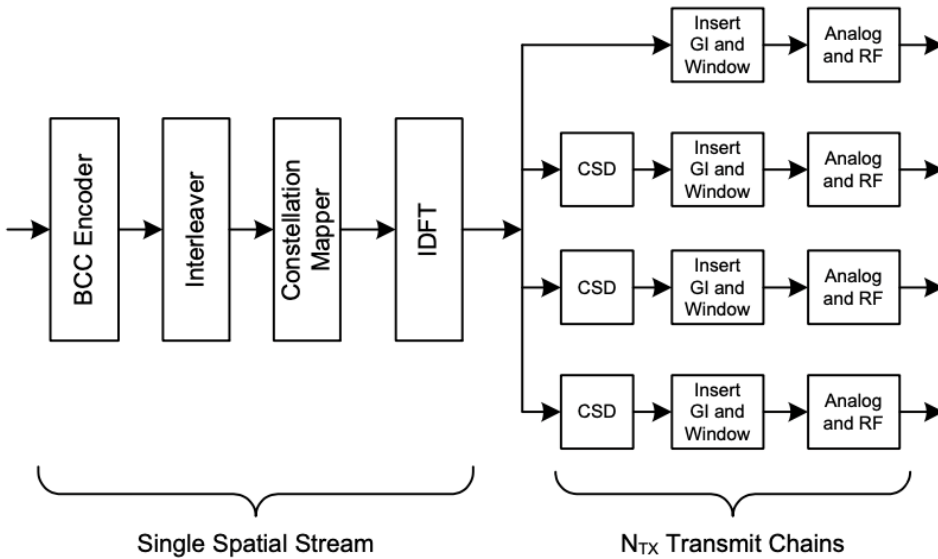


Figure 20-2—Transmitter block diagram 1

Source: IEEE Std. 802.11-2012, p. 1684-85.

20.3.3 Transmitter block diagram

HT-mixed format and HT-greenfield format transmissions can be generated using a transmitter consisting of the following blocks:

- a) *Scrambler* scrambles the data to reduce the probability of long sequences of 0s or 1s; see 20.3.11.3.
- b) *Encoder parser*, if BCC encoding is to be used, demultiplexes the scrambled bits among N_{ES} (number of BCC encoders for the Data field) BCC encoders, in a round robin manner.
- c) *FEC encoders* encode the data to enable error correction. An FEC encoder may include a binary convolutional encoder followed by a puncturing device, or it may include an LDPC encoder.
- d) *Stream parser* divides the outputs of the encoders into blocks that are sent to different interleaver and mapping devices. The sequence of the bits sent to an interleaver is called a *spatial stream*.
- e) *Interleaver* interleaves the bits of each spatial stream (changes order of bits) to prevent long sequences of adjacent noisy bits from entering the BCC decoder. Interleaving is applied only when BCC encoding is used.
- f) *Constellation mapper* maps the sequence of bits in each spatial stream to constellation points (complex numbers).
- g) *STBC encoder* spreads constellation points from N_{SS} spatial streams into N_{STS} space-time streams using a space-time block code. STBC is used only when $N_{SS} < N_{STS}$; see 20.3.11.9.2.

- h) *Spatial mapper* maps space-time streams to transmit chains. This may include one of the following:
 - 1) *Direct mapping*: Constellation points from each space-time stream are mapped directly onto the transmit chains (one-to-one mapping).
 - 2) *Spatial expansion*: Vectors of constellation points from all the space-time streams are expanded via matrix multiplication to produce the input to all the transmit chains.
 - 3) *Beamforming*: Similar to spatial expansion, each vector of constellation points from all the space-time streams is multiplied by a matrix of steering vectors to produce the input to the transmit chains.
- i) *Inverse discrete Fourier transform (IDFT)* converts a block of constellation points to a time domain block.

Source: IEEE Std. 802.11-2012, p. 1683-84.

133. The method practiced by the '439 Accused Products includes cyclically advancing the first OFDM packet by shifting the samples in a first direction an amount less than a sample duration of the guard interval portion to generate a shifted version of the first OFDM packet for transmission. For example, the Accused Products cyclically shift the symbol data portion of the HT_SIG by -200 ns up to -50 ns, which is less than its guard interval's total length of 0.8 us, to generate a shifted version for transmission as seen below:

20.3.3 Transmitter block diagram

HT-mixed format and HT-greenfield format transmissions can be generated using a transmitter consisting of the following blocks:

- a) *Scrambler* scrambles the data to reduce the probability of long sequences of 0s or 1s; see 20.3.11.3.
- b) *Encoder parser*, if BCC encoding is to be used, demultiplexes the scrambled bits among N_{ES} (number of BCC encoders for the Data field) BCC encoders, in a round robin manner.
- c) *FEC encoders* encode the data to enable error correction. An FEC encoder may include a binary convolutional encoder followed by a puncturing device, or it may include an LDPC encoder.
- d) *Stream parser* divides the outputs of the encoders into blocks that are sent to different interleaver and mapping devices. The sequence of the bits sent to an interleaver is called a *spatial stream*.
- e) *Interleaver* interleaves the bits of each spatial stream (changes order of bits) to prevent long sequences of adjacent noisy bits from entering the BCC decoder. Interleaving is applied only when BCC encoding is used.
- f) *Constellation mapper* maps the sequence of bits in each spatial stream to constellation points (complex numbers).
- g) *STBC* encoder spreads constellation points from N_{SS} spatial streams into N_{STS} space-time streams using a space-time block code. STBC is used only when $N_{SS} < N_{STS}$; see 20.3.11.9.2.

- h) *Spatial mapper* maps space-time streams to transmit chains. This may include one of the following:
- 1) *Direct mapping*: Constellation points from each space-time stream are mapped directly onto the transmit chains (one-to-one mapping).
 - 2) *Spatial expansion*: Vectors of constellation points from all the space-time streams are expanded via matrix multiplication to produce the input to all the transmit chains.
 - 3) *Beamforming*: Similar to spatial expansion, each vector of constellation points from all the space-time streams is multiplied by a matrix of steering vectors to produce the input to the transmit chains.
- i) *Inverse discrete Fourier transform (IDFT)* converts a block of constellation points to a time domain block.

Source: IEEE Std. 802.11-2012, p. 1683-84.




Figure 20-2—Transmitter block diagram 1

Source: IEEE Std. 802.11-2012, p. 1685.

20.3.9.3.2 Cyclic shift definition

The cyclic shift values defined in this subclause apply to the non-HT fields in the HT-mixed format preamble and the HT-SIG in the HT-mixed format preamble.

Cyclic shifts are used to prevent unintentional beamforming when the same signal or scalar multiples of one signal are transmitted through different spatial streams or transmit chains. A cyclic shift of duration T_{CS} on a signal $s(t)$ on interval $0 \leq t \leq T$ is defined as follows, where T is defined as T_{DFT} as referenced in Table 20-6.

With $T_{CS} \leq 0$, replace $s(t)$ with $s(t - T_{CS})$ when $0 \leq t < T + T_{CS}$ and with $s(t - T_{CS} - T)$ when $T + T_{CS} \leq t \leq T$. The cyclic-shifted signal is defined as shown in Equation (20-7).

$$s_{CS}(t; T_{CS}) \Big|_{T_{CS} < 0} = \begin{cases} s(t - T_{CS}) & 0 \leq t < T + T_{CS} \\ s(t - T_{CS} - T) & T + T_{CS} \leq t \leq T \end{cases} \quad (20-7)$$

The cyclic shift is applied to each OFDM symbol in the packet separately. Table 20-9 specifies the values for the cyclic shifts that are applied in the L-STF (in an HT-mixed format packet), the L-LTF, and L-SIG. It also applies to the HT-SIG in an HT-mixed format packet.

Source: IEEE Std. 802.11-2012, p. 1694-95.

Table 20-9—Cyclic shift for non-HT portion of packet

T_{CS}^{TX} values for non-HT portion of packet				
Number of transmit chains	Cyclic shift for transmit chain 1 (ns)	Cyclic shift for transmit chain 2 (ns)	Cyclic shift for transmit chain 3 (ns)	Cyclic shift for transmit chain 4 (ns)
1	0	—	—	—
2	0	-200	—	—
3	0	-100	-200	—
4	0	-50	-100	-150

Source: IEEE Std. 802.11-2012, p. 1695.

Table 20-5—Timing-related constants (continued)

T_{DFT} : IDFT/DFT period	3.2 μ s
T_{GI} : Guard interval duration	0.8 μ s = $T_{DFT}/4$

Source: IEEE Std. 802.11n-2009, p. 266.

20.3.4 Overview of the PPDU encoding process

The encoding process is composed of the steps described below. The following overview is intended to facilitate an understanding of the details of the convergence procedure:

- b) Construct the PLCP preamble SIGNAL fields from the appropriate fields of the TXVECTOR by adding tail bits, applying convolutional coding, formatting into one or more OFDM symbols, applying cyclic shifts, applying spatial processing, calculating an inverse Fourier transform for each OFDM symbol and transmit chain, and prepending a cyclic prefix or GI to each OFDM symbol in each transmit chain. The number and placement of the PLCP preamble SIGNAL fields depend on the frame format being used. Refer to 20.3.9.3.5, 20.3.9.4.3, and 20.3.9.5.4.
- r) For each group of N_{ST} subcarriers and each of the N_{TX} transmit chains, convert the subcarriers to time domain using IDFT. Prepend to the Fourier-transformed waveform a circular extension of itself, thus forming a GI, and truncate the resulting periodic waveform to a single OFDM symbol length by applying time domain windowing. Determine the length of the GI according to the GI_TYPE parameter of the TXVECTOR. Refer to 20.3.11.11 and 20.3.11.12 for details. When beamforming is not used, it is sometimes possible to implement the cyclic shifts in the time domain.

Source: IEEE Std. 802.11-2012, p. 1684, 1688.


The Cyclic Prefix for OFDM

In a previous post, we have elaborated about the [building blocks of OFDM](#).

There, we have stated two benefits of using a cyclic prefix between subsequent OFDM symbols:

- The CP isolates different OFDM blocks from each other when the wireless channel contains multiple paths, i.e. is frequency-selective.
- The CP turns the linear convolution with the channel into a [circular convolution](#). Only with a circular convolution, we can use the single-tap equalization OFDM is so famous for.

As we see, the CP of an OFDM symbol is obtained by prepending a copy of the last N_{CP} samples from the end of the OFDM signal to its beginning. This way we obtain a circular signal structure, i.e. the first N_{CP} and last N_{CP} samples are equal in each OFDM symbol.



In the above figure, we see two subsequent OFDM symbols, each having a dedicated CP. The colors encode the signal value. The cyclic prefix at the beginning of each OFDM symbol shows a copy of the color of end of the OFDM symbol. When the signal is demodulated, the N-point FFT is taken at the position after the CP, which is indicated with *FFT window*.

Source: <https://dspillustrations.com/pages/posts/misc/the-cyclic-prefix-cp-in-ofdm.html>

The HT-SIG is composed of two parts, HT-SIG₁ and HT-SIG₂, each containing 24 bits, as shown in Figure 20-6. All the fields in the HT-SIG are transmitted LSB first, and HT-SIG₁ is transmitted before HT-SIG₂.

The HT-SIG parts shall be encoded at $R = 1/2$, interleaved, and mapped to a BPSK constellation, and they have pilots inserted following the steps described in 18.3.5.6, 18.3.5.7, 18.3.5.8, and 18.3.5.9, respectively. The BPSK constellation is rotated by 90° relative to the L-SIG in order to accommodate detection of the start of the HT-SIG. The stream of 96 complex numbers generated by these steps is divided into two groups of 48 complex numbers: $d_{k,n}$, $0 \leq k \leq 47$, $n = 0, 1$. The time domain waveform for the HT-SIG in an HT-mixed format packet in a 20 MHz transmission shall be as shown in Equation (20-16).

$$r_{HT-SIG}^{i_{rx}}(t) = \frac{1}{\sqrt{N_{TX} \cdot N_{HT-SIG}^{Tone}}} \sum_{n=0}^1 w_{T_{SYM}}(t - nT_{SYM}) \cdot \sum_{k=-26}^{26} (jD_{k,n} + p_{n+1}P_k) \exp(j2\pi k\Delta_F(t - nT_{SYM} - T_{GI} - T_{CS}^{i_{rx}})) \quad (20-16)$$

where

$$D_{k,n} = \begin{cases} 0, & k = 0, \pm 7, \pm 21 \\ d_{M'(k),n} & \text{otherwise} \end{cases}$$

$M'(k)$ is defined in 20.3.9.3

P_k and p_n are defined in 18.3.5.10

N_{HT-SIG}^{Tone} has the value given in Table 20-8

$T_{CS}^{i_{TX}}$ represents the cyclic shift for transmit chain i_{TX} and is defined by Table 20-9 for HT-mixed format PPDU.

Source: IEEE Std. 802.11-2012, pp. 1700-1701.

T_{GI} : Double guard interval	1.6 μ s	1.6 μ s	1.6 μ s
----------------------------------	-------------	-------------	-------------

Source: IEEE Std. 802.11-2012, p. 1689.

134. In the method practiced by the '439 Accused Products, a number of samples from the symbol data portion of the first OFDM packet are shifted into the guard interval portion of the shifted version of the first OFDM packet, and the same number of samples from the guard interval portion of the first OFDM packet are shifted out of the guard interval portion of the shifted version of the first OFDM packet. For example, the Accused Products cyclically advance the number of samples corresponding to the time duration of $/T_{CS}/$ out of the symbol portion of the shifted OFDM packet and into the guard interval portion of the packet, while the same number of samples corresponding to the time duration of $/T_{CS}/$ are shifted out of the guard interval portion of the shifted OFDM packet, as illustrated below:

20.3.3 Transmitter block diagram

HT-mixed format and HT-greenfield format transmissions can be generated using a transmitter consisting of the following blocks:

- a) *Scrambler* scrambles the data to reduce the probability of long sequences of 0s or 1s; see 20.3.11.3.
- b) *Encoder parser*, if BCC encoding is to be used, demultiplexes the scrambled bits among N_{ES} (number of BCC encoders for the Data field) BCC encoders, in a round robin manner.
- c) *FEC encoders* encode the data to enable error correction. An FEC encoder may include a binary convolutional encoder followed by a puncturing device, or it may include an LDPC encoder.
- d) *Stream parser* divides the outputs of the encoders into blocks that are sent to different interleaver and mapping devices. The sequence of the bits sent to an interleaver is called a *spatial stream*.
- e) *Interleaver* interleaves the bits of each spatial stream (changes order of bits) to prevent long sequences of adjacent noisy bits from entering the BCC decoder. Interleaving is applied only when BCC encoding is used.
- f) *Constellation mapper* maps the sequence of bits in each spatial stream to constellation points (complex numbers).
- g) *STBC encoder* spreads constellation points from N_{SS} spatial streams into N_{STS} space-time streams using a space-time block code. STBC is used only when $N_{SS} < N_{STS}$; see 20.3.11.9.2.
- h) *Spatial mapper* maps space-time streams to transmit chains. This may include one of the following:
 - 1) *Direct mapping*: Constellation points from each space-time stream are mapped directly onto the transmit chains (one-to-one mapping).
 - 2) *Spatial expansion*: Vectors of constellation points from all the space-time streams are expanded via matrix multiplication to produce the input to all the transmit chains.
 - 3) *Beamforming*: Similar to spatial expansion, each vector of constellation points from all the space-time streams is multiplied by a matrix of steering vectors to produce the input to the transmit chains.
- i) *Inverse discrete Fourier transform (IDFT)* converts a block of constellation points to a time domain block.

Source: IEEE Std. 802.11-2012, p. 1683-84.




Figure 20-2—Transmitter block diagram 1

Source: IEEE Std. 802.11-2012, p. 1685.

20.3.9.3.2 Cyclic shift definition

The cyclic shift values defined in this subclause apply to the non-HT fields in the HT-mixed format preamble and the HT-SIG in the HT-mixed format preamble.

Cyclic shifts are used to prevent unintentional beamforming when the same signal or scalar multiples of one signal are transmitted through different spatial streams or transmit chains. A cyclic shift of duration T_{CS} on a signal $s(t)$ on interval $0 \leq t \leq T$ is defined as follows, where T is defined as T_{DFT} as referenced in Table 20-6.

With $T_{CS} \leq 0$, replace $s(t)$ with $s(t - T_{CS})$ when $0 \leq t < T + T_{CS}$ and with $s(t - T_{CS} - T)$ when $T + T_{CS} \leq t \leq T$. The cyclic-shifted signal is defined as shown in Equation (20-7).

$$s_{CS}(t; T_{CS})|_{T_{CS} < 0} = \begin{cases} s(t - T_{CS}) & 0 \leq t < T + T_{CS} \\ s(t - T_{CS} - T) & T + T_{CS} \leq t \leq T \end{cases} \quad (20-7)$$

The cyclic shift is applied to each OFDM symbol in the packet separately. Table 20-9 specifies the values for the cyclic shifts that are applied in the L-STF (in an HT-mixed format packet), the L-LTF, and L-SIG. It also applies to the HT-SIG in an HT-mixed format packet.

Source: IEEE Std. 802.11-2012, p. 1694-95.

Table 20-9—Cyclic shift for non-HT portion of packet

T_{CS}^{fix} values for non-HT portion of packet				
Number of transmit chains	Cyclic shift for transmit chain 1 (ns)	Cyclic shift for transmit chain 2 (ns)	Cyclic shift for transmit chain 3 (ns)	Cyclic shift for transmit chain 4 (ns)
1	0	—	—	—
2	0	-200	—	—
3	0	-100	-200	—
4	0	-50	-100	-150

Source: IEEE Std. 802.11-2012, p. 1695.

Table 20-5—Timing-related constants (continued)

T_{DFT} : IDFT/DFT period	3.2 μ s
T_{GI} : Guard interval duration	0.8 μ s = $T_{DFT}/4$

Source: IEEE Std. 802.11n-2009, p. 266.

20.3.4 Overview of the PPDU encoding process

The encoding process is composed of the steps described below. The following overview is intended to facilitate an understanding of the details of the convergence procedure:

- b) Construct the PLCP preamble SIGNAL fields from the appropriate fields of the TXVECTOR by adding tail bits, applying convolutional coding, formatting into one or more OFDM symbols, applying cyclic shifts, applying spatial processing, calculating an inverse Fourier transform for each OFDM symbol and transmit chain, and prepending a cyclic prefix or GI to each OFDM symbol in each transmit chain. The number and placement of the PLCP preamble SIGNAL fields depend on the frame format being used. Refer to 20.3.9.3.5, 20.3.9.4.3, and 20.3.9.5.4.

- r) For each group of N_{ST} subcarriers and each of the N_{TX} transmit chains, convert the subcarriers to time domain using IDFT. Prepend to the Fourier-transformed waveform a circular extension of itself, thus forming a GI, and truncate the resulting periodic waveform to a single OFDM symbol length by applying time domain windowing. Determine the length of the GI according to the GI_TYPE parameter of the TXVECTOR. Refer to 20.3.11.11 and 20.3.11.12 for details. When beamforming is not used, it is sometimes possible to implement the cyclic shifts in the time domain.

Source: IEEE Std. 802.11-2012, p. 1684, 1688.


The Cyclic Prefix for OFDM

In a previous post, we have elaborated about the [building blocks of OFDM](#).

There, we have stated two benefits of using a cyclic prefix between subsequent OFDM symbols:

- The CP isolates different OFDM blocks from each other when the wireless channel contains multiple paths, i.e. is frequency-selective.
- The CP turns the linear convolution with the channel into a [circular convolution](#). Only with a circular convolution, we can use the single-tap equalization OFDM is so famous for.

As we see, the CP of an OFDM symbol is obtained by prepending a copy of the last N_{CP} samples from the end of the OFDM signal to its beginning. This way we obtain a circular signal structure, i.e. the first N_{CP} and last N_{CP} samples are equal in each OFDM symbol.



In the above figure, we see two subsequent OFDM symbols, each having a dedicated CP. The colors encode the signal value. The cyclic prefix at the beginning of each OFDM symbol shows a copy of the color of end of the OFDM symbol. When the signal is demodulated, the N -point FFT is taken at the position after the CP, which is indicated with *FFT window*.

Source: <https://dspillustrations.com/pages/posts/misc/the-cyclic-prefix-cp-in-ofdm.html>

The HT-SIG is composed of two parts, HT-SIG₁ and HT-SIG₂, each containing 24 bits, as shown in Figure 20-6. All the fields in the HT-SIG are transmitted LSB first, and HT-SIG₁ is transmitted before HT-SIG₂.

The HT-SIG parts shall be encoded at $R = 1/2$, interleaved, and mapped to a BPSK constellation, and they have pilots inserted following the steps described in 18.3.5.6, 18.3.5.7, 18.3.5.8, and 18.3.5.9, respectively. The BPSK constellation is rotated by 90° relative to the L-SIG in order to accommodate detection of the start of the HT-SIG. The stream of 96 complex numbers generated by these steps is divided into two groups of 48 complex numbers: $d_{k,n}$, $0 \leq k \leq 47$, $n = 0, 1$. The time domain waveform for the HT-SIG in an HT-mixed format packet in a 20 MHz transmission shall be as shown in Equation (20-16).

$$r_{HT-SIG}^{i_{TX}}(t) = \frac{1}{\sqrt{N_{TX} \cdot N_{HT-SIG}^{Tone}}} \sum_{n=0}^1 w_{T_{SYM}}(t - nT_{SYM}) \cdot \sum_{k=-26}^{26} (jD_{k,n} + p_{n+1}P_k) \exp(j2\pi k\Delta_F(t - nT_{SYM} - T_{GI} - T_{CS}^{i_{TX}})) \quad (20-16)$$

where

$$D_{k,n} = \begin{cases} 0, & k = 0, \pm 7, \pm 21 \\ d_{M'(k),n}, & \text{otherwise} \end{cases}$$

$M'(k)$ is defined in 20.3.9.3

P_k and p_n are defined in 18.3.5.10

N_{HT-SIG}^{Tone} has the value given in Table 20-8

$T_{CS}^{i_{TX}}$ represents the cyclic shift for transmit chain i_{TX} and is defined by Table 20-9 for HT-mixed format PPDU.

Source: IEEE Std. 802.11-2012, pp. 1700-1701.

T_{GI} : Double guard interval	1.6 μ s	1.6 μ s	1.6 μ s
----------------------------------	-------------	-------------	-------------

Source: IEEE Std. 802.11-2012, p. 1689.

2.7.4 Mixed Mode Preamble

The mixed mode preamble is needed for compatibility with IEEE 802.11a/g. It starts with the 802.11a/g preamble. The 802.11n mixed mode preamble for two spatial streams is shown in Figure 2.34. The legacy short training field (L-STF) is identical to 802.11a/g except that different transmitters use different cyclic delays (CDs). This also applies to the legacy long training field (L-LTF). The STFs from different transmitters have low cross-correlation. For example, a CD of -400 ns (or a cyclic advance of 400 ns) minimizes correlation between two different transmitted short symbols. The L-STF uses a CD of only -200 ns for two transmitters, since legacy 802.11a/g receivers may not be able to cope with larger CD values.

Source: B. Bing, Broadband Wireless Multimedia Networks, Wiley, 2013, p. 120.

With $T_{CS} \leq 0$, replace $s(t)$ with $s(t - T_{CS})$ when $0 \leq t < T + T_{CS}$ and with $s(t - T_{CS} - T)$ when $T + T_{CS} \leq t \leq T$. The cyclic-shifted signal is defined as shown in Equation (20-7).

$$s_{CS}(t; T_{CS})|_{T_{CS} < 0} = \begin{cases} s(t - T_{CS}) & 0 \leq t < T + T_{CS} \\ s(t - T_{CS} - T) & T + T_{CS} \leq t \leq T \end{cases} \quad (20-7)$$

135. The method practiced by the '439 Accused Products includes substantially simultaneously transmitting the first OFDM packet and the shifted version of the OFDM packet. For example, the signals transmitted from different transmit chains are aligned and synchronized in the time domain, as seen below:

Figure 20-2 and Figure 20-3 show example transmitter block diagrams. In particular, Figure 20-2 shows the transmitter blocks used to generate the HT-SIG of the HT-mixed format PPDU.

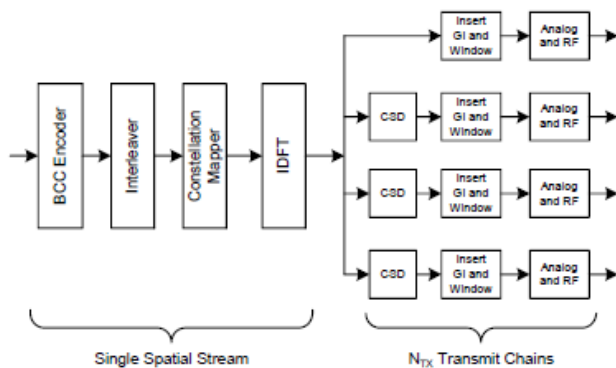


Figure 20-2—Transmitter block diagram 1

Source: IEEE Std. 802.11-2012, pp. 1684-85.

Radio Chains

Between the operating system and antenna, an 802.11 radio interface has to perform several tasks. When transmitting a frame, the main tasks are the inverse Fourier transform to turn the frequency-domain encoded signal into a time-domain signal, and amplification right before the signal hits the antenna so it has reasonable range. On the receive side, the process must be reversed. Immediately after entering the antenna, an amplifier boosts the faint signal received into something substantial enough to work with, and performs a Fourier transform to extract the subcarriers. In an 802.11 interface, these components are linked together and called a *radio chain*. Selecting the components to make up the radio chain is an important task for system designers, especially

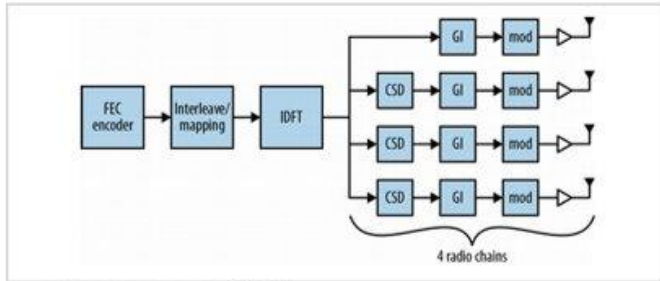


Figure 2-3. 4x4 802.11n interface block diagram

Source: Gast, Matthew, S., 802.11n A Survival Guide, O'Reilly, 2012, pp. 13-14.

136. Additionally, Defendant has been and currently is an active inducer of infringement of the '439 patent under 35 U.S.C. § 271(b) and a contributory infringer of the '439 patent under 35 U.S.C. § 271(c).

137. Defendant has actively induced, and continues to actively induce, infringement of the '439 patent by causing others to use, offer for sale, or sell in the United States, products or services covered by the '439 patent, including but not limited to the '439 Accused Products and any other products or services that include WiFi chipsets compliant with 802.11n, 802.11ac and/or 802.11ax, having the cyclic shift advance functionality described above. Defendant provides these products and services to others, such as customers, resellers, partners, and end-users, who, in turn, use, provision for use, offer for sale, or sell those products and services, which directly infringe the '439 patent as described above. Defendant's inducement includes requiring WiFi chipsets within the Accused Products to be compliant with the IEEE 802.11n, 802.11ac and 802.11ax standard, in which the cyclic advance diversity scheme described above is mandatory, and advertising and promoting such compliance to its customers, partners, re-sellers and the like, including the promotion, directions and instructions found at one or more of the following links, the provision of which has been on-going since the filing of the First Amended Complaint in case 6:23-cv-0068 and the content of which is specifically illustrated above:

- <https://support.lenovo.com/us/en/solutions/ht070352-introduction-to-some-technical-parameters-of-lenovo-80211n-wireless-adapter>
- <https://support.lenovo.com/us/en/solutions/HF001441>
- <https://forums.lenovo.com/t5/Moto-Z3-Play/Moto-Z3-Wi-fi-won-t-connect-at-300-Mbps-with-older-routers-while-Z2-does/m-p/4288541>
- https://www.wi-fi.org/product-finder-results?sort_by=certified&sort_order=desc&certifications=276&keywords=motorola
- https://www.wi-fi.org/product-finder-results?sort_by=certified&sort_order=desc&certifications=276&keywords=lenovo&companies=466
- <https://thinkstation-specs.com/wp-content/uploads/2021/06/P15-Gen-2-Lenovo-ThinkStation.pdf>
- <https://www.lenovo.com/us/en/p/phones/motorola-smartphones/motorola-razr/pajs0007us>
- [https://www.lenovo.com/us/en/p/phones/motorola-smartphones/motorola-moto-g/moto-g-stylus-5g-\(2022\)/patj0009us](https://www.lenovo.com/us/en/p/phones/motorola-smartphones/motorola-moto-g/moto-g-stylus-5g-(2022)/patj0009us)
- [https://www.lenovo.com/us/en/p/phones/motorola-smartphones/motorola-moto-g/moto-g-power-\(2022\)/pase0012us](https://www.lenovo.com/us/en/p/phones/motorola-smartphones/motorola-moto-g/moto-g-power-(2022)/pase0012us)
- [https://www.lenovo.com/us/en/p/phones/motorola-smartphones/motorola-moto-g/moto-g-stylus-\(2022\)/pat40000us](https://www.lenovo.com/us/en/p/phones/motorola-smartphones/motorola-moto-g/moto-g-stylus-(2022)/pat40000us)
- [https://www.lenovo.com/us/en/p/phones/motorola-smartphones/motorola-moto-g/moto-g-power-\(2022\)/pase0005us](https://www.lenovo.com/us/en/p/phones/motorola-smartphones/motorola-moto-g/moto-g-power-(2022)/pase0005us)
- <https://www.motorola.com/us/smartphones-motorola-edge-plus/p>

138. Defendant has contributed to, and continues to contribute to, the infringement of the '439 patent by others by knowingly providing one or more components, for example the 802.11 WiFi chipset with cyclic shift (advance) functionality included in the Accused Products, a portion thereof, and/or the software/hardware modules responsible for the accused functionality described herein, that, when installed, configured, and used result in systems that, as intended by Defendant described above, directly infringe one or more claims of the '439 patent.

139. Defendant knew of the '439 patent, or should have known of the '439 patent, but was willfully blind to its existence. Upon information and belief, Defendant had actual knowledge of the '439 patent since at least as early as the filing of the First Amended Complaint in case 6:23-cv-0068-ADA, or alternatively, at least as early as Defendant's receipt of this Complaint.

Alternatively, upon information and belief, Defendant has had knowledge of the '439 patent since the service upon Defendant of the Complaint in this action.

140. By the time of trial, Defendant will or should have known and intended (since receiving such notice) that its continued actions would infringe and would actively induce and contribute to the infringement of the '439 patent.

141. Defendant has committed, and continues to commit, contributory infringement by selling products and services that directly infringe the '439 patent when used by a third party, such as the Accused '439 Products, and that are a material part of the invention, knowing them to be especially made or adapted for use in infringement of the '439 patent and not staple articles or commodities of commerce suitable for substantial non-infringing use.

142. As a result of Defendant's acts of infringement, IV has suffered and will continue to suffer damages in an amount to be determined at trial.

COUNT V

(Defendant's Infringement of U.S. Patent No. 7,646,835).

143. The preceding paragraphs are reincorporated by reference as if fully set forth herein.

144. The '835 patent claims and teaches, *inter alia*, an improved way to input and output signaling for digital integrated circuit devices, by way of automatically calibrating intra-cycle timing relationships between command signals, data signals, and sampling signals for an integrated circuit device.

145. The inventions improved upon the then-existing manual calibration techniques that were used on high-speed integrated circuit device memory systems, and other types of high-performance integrated circuit devices that require precisely aligned signals for their input and output. As was well-known, the design and certification of high-speed DDR memory modules

had become a significant challenge for many system manufacturers. Practically all of the integral features of a given DDR dual in-line memory module (“DIMM”), such as the particular type of silicon used, the routing and thickness of the printed circuit board (“PCB”), and the signal integrity performance under stressed conditions (e.g., temperature, voltage, etc.), have an impact on the overall system performance and reliability. Failure to properly account for these variables can result in single or multi-bit errors, read/write command sequencing failures, failure of the system to attain rated performance, and the like.

146. The automatic calibration process as provided by the inventions of the '835 patent add a significant amount of “extra margin” to the specifications of a memory system. For example, for integrated circuit devices such as DDR DRAMs, the DDR timing specifications are so stringent that even slight variations (e.g., between motherboards, devices from different lots, etc.) can have an impact on overall system performance and cause intermittent timing-related DIMM failures. These failures can be the most difficult types of failures to detect and correct. The extra margin provided by the automatic calibration inventions of the '835 patent increase the reliability rate of computer systems incorporating such high-performance integrated circuit devices. The extra margin provided by embodiments of the '835 patent can be used to increase the maximum obtainable performance of such computer systems.

147. More specifically, the claims of the '835 patent recite automatic calibration of intra-cycle timing relationships between command signals, data signals, and sampling signals for an integrated circuit device, including the generation of command signals to access an integrated circuit component; the accessing of data signals to convey data for the integrated circuit component; the accessing of sampling signals to control sampling of the data signals; and systematically altering a phase shift of the command signals, a phase shift of the data signals, and

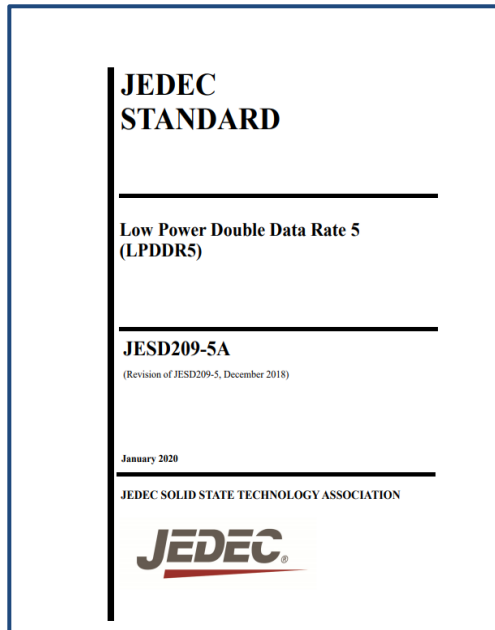
a phase shift of the sampling signals to determine a valid operation range of the integrated circuit device, wherein the valid operation range includes an optimal operation point for the integrated circuit device.

148. The method covered by the asserted claims, therefore, differs markedly from the prior systems in use at the time of this invention, which lacked the claimed combination of automatic calibration of intra-cycle timing relationships between command signals, data signals, and sampling signals for an integrated circuit device, including the generation of command signals to access an integrated circuit component; the accessing of data signals to convey data for the integrated circuit component; the accessing of sampling signals to control sampling of the data signals; and systematically altering a phase shift of the command signals, a phase shift of the data signals, and a phase shift of the sampling signals to determine a valid operation range of the integrated circuit device, wherein the valid operation range includes an optimal operation point for the integrated circuit device.

149. Defendant has directly infringed and continues to directly infringe at least claim 1 of the '835 patent by making, using, testing, selling, offering for sale, and importing into the United States products and services covered by the '835 patent. Defendant's products and services that infringe the '835 patent include all mobile phones, laptops and tablets that include Qualcomm-based processors and LPDDR4, LPDDR4X or LPDDR5 memory, (together the "Accused '835 Products" or "Accused Products"). Specific examples of the Accused Products include, but are not limited to, the Lenovo 10W-10.1 Snapdragon tablet, Motorola Edge+, and Lenovo ThinkPad S13s Snapdragon laptop.

150. LPDDR4, LPDDR4X and LPDDR5 memory (including that within the Accused Products), are required to follow the JEDEC Solid State Technology Association's standard for

Low Power Double Data Rate 4, 4X and 5. Upon information and belief, memory labeled with the designations LPDDR4, LPDDR4X and LPDDR5 are claiming conformance with the standard by using such designations. The JEDEC standards for LPDDR4, LPDDR4X and LPDDR5 operate substantially the same with respect to the accused functionality (as seen below), and therefore, the LPDDR5 standard will be used as exemplary.



1 Scope

This document defines the LPDDR5 standard, including features, functionalities, AC and DC characteristics, packages, and ball/signal assignments. The purpose of this specification is to define the minimum set of requirements for a JEDEC compliant x16 one channel SDRAM device and x8 one channel SDRAM device. LPDDR5 device density ranges from 2 Gb through 32 Gb. This document was created using aspects of the following standards: DDR2 (JESD79-2), DDR3 (JESD79-3), DDR4 (JESD79-4), LPDDR (JESD209), LPDDR2 (JESD209-2), LPDDR3 (JESD209-3) and LPDDR4 (JESD209-4).

Each aspect of the standard was considered and approved by committee ballot(s). The accumulation of these ballots was then incorporated to prepare the LPDDR5 standard.

No claims to be in conformance with this standard may be made unless all requirements stated in the standard are met.

Source: JEDEC Standard No. 209-5A - Cover Page; p.1 and Notice Page.

I/O Signal Trainings

There are multiple trainings provided by LPDDR4 to align or re-adjust the delays introduced on the I/O signals with respect to CLK or other signals. As per standard physical interface definition of LPDDR4, there are CLK, CS, CA, DQ and DQS signals which need proper alignment for successful data transfers. As the CA line is sampled at the CLK signal, there should be a proper phase relationship between CA and CLK. Similarly, DQ gets sampled on DQS signal, so again there should be a phase relationship between the two. To maintain these phase relationships, LPDDR4 proposes training mechanisms. Let's look at those:

- **Command Bus Training (CBT):** This is used to align the CS and CA signals with respect to the CLK signal. At power-up the receivers get configured for low speed operations. When operating at high frequencies, the receivers must be trained. The timing margins need to be readjusted per the higher clock frequency which is achieved with the CBT procedure. The entry and exit of the CBT mode are controlled by the mode register write command. In CBT mode, DRAM will switch to the FSP_OP settings, which it will also need to be trained on. DRAM samples the CA bus at CS signal and provides feedback of the sampled signals to the controller for timing adjustments on CS and CA signals.
- **Write Leveling:** This is used to adjust the delays on DQS input signals with respect to the CLK signal. The entry and exit of the write leveling training mode are controlled by the mode register write command. DQS signal gets driven by the controller and DRAM samples the CLK signal at the DQS edge. DRAM responds to the controller by providing feedback on the captured CLK level on DQ. This feedback identifies the leading or lagging of DQS, with respect to the CLK, so that controller can readjust the delays accordingly.
- **Write Training (DQS-DQ Training):** This is used to align the DQ input signal delays with respect to the DQS input signal. When entering write training mode, MPC WR_DQ_FIFO command must be issued by the controller. This command writes a user defined data in DRAM, then the controller issues MPC RD_DQ_FIFO command to read back the data from the same location and compare both the written and read data to re-adjust the delay on DQ line.

Source: <https://blogs.synopsys.com/vip-central/2017/10/03/lpddr4-the-total-package-for-mobile-soc-ram/>

Low Power Double Data Rate 5 JESD209-5A (LPDDR5)

(Revision of JESD209-5, December 2018)

4.2.2 Command Bus Training

The LPDDR5 SDRAM command bus must be trained before enabling termination for high-frequency or mid-frequency operation. LPDDR5 SDRAM provides an internal VREF(CA) that default level is suitable for un-terminated, low-frequency operation, however the VREF(CA) must be trained to achieve suitable receiver voltage margin for terminated, high-frequency or mid-frequency operation. The training mode described here centers the internal VREF(CA) in the CA data eye and at the same time allows for timing adjustments of the CS and CA signals to meet their Rx Mask requirements. For the training sequence simplicity and difficulty to capture CA inputs prior to training the CA inputs, the training mode described here uses a minimum of external commands to enter, train, and exit the Command Bus Training.

4.2.5 WCK2CK Leveling

4.2.5.1 WCK2CK Leveling Mode (write-leveling called in LPDDR4)

To adjust CK-to-WCK relationship and guarantee WCK2CK-Sync. operation, the LPDDR5 SDRAM provides a WCK2CK Leveling feature to compensate CK-to-WCK timing skew affecting WCK2CK-Sync. operation. The SDRAM compares the phase of the rising edge of WCK and the rising edge of CK, then asynchronously feeds back to the memory controller for the WCK2CK phase detection result. After finishing WCK2CK Leveling, tWCK2CK which means CK-to-WCK relationship is determined and WCK2CK-Sync. operation will be performed with the optimized margin.

4.2.9 WCK-DQ Training

The LPDDR5 SDRAM uses an un-matched WCK-DQ path to enable high speed performance and save power in the SDRAM. As a result, WCK is required to be trained to arrive at the DQ latch center-aligned with the Data eye. The SDRAM DQ receiver is located at the DQ pad and has a shorter internal delay in the SDRAM than does the WCK signal. The SDRAM DQ receiver will latch the data present on the DQ bus when WCK reaches the latch, and training is accomplished by delaying the DQ signals relative to WCK such that the Data eye arrives at the receiver latch centered on the WCK transition.

Source: JEDEC Standard No. 209-5A at 46, 75 & 90.

151. Claim 1 of the '835 patent is reproduced below:

1. A method for automatically calibrating intra-cycle timing relationships between command signals, data signals, and sampling signals for an integrated circuit device, the method comprising:

generating command signals to access an integrated circuit component;

accessing data signals to convey data for the integrated circuit component;

accessing sampling signals to control sampling of the data signals;
and

systematically altering a phase shift of the command signals, a phase shift of the data signals, and a phase shift of the sampling signals to determine a valid operation range of the integrated circuit device, wherein the valid operation range includes an optimal operation point for the integrated circuit device.

152. The Accused '835 Products each automatically calibrate intra-cycle timing relationships between command signals, data signals, and sampling signals for an integrated circuit device. For example, the Qualcomm Snapdragon 8 Gen 1 Mobile Platform requires LPDDR5 training of command, clocking and data symbols during initialization to establish critical timing relationships prior to normal operation, as seen below:

Low Power Double Data Rate 5 **JESD209-5A** (LPDDR5)

(Revision of JESD209-5, December 2018)

4.2.2 Command Bus Training

The LPDDR5 SDRAM command bus must be trained before enabling termination for high-frequency or mid-frequency operation. LPDDR5 SDRAM provides an internal VREF(CA) that default level is suitable for un-terminated, low-frequency operation, however the VREF(CA) must be trained to achieve suitable receiver voltage margin for terminated, high-frequency or mid-frequency operation. The training mode described here centers the internal VREF(CA) in the CA data eye and at the same time allows for timing adjustments of the CS and CA signals to meet their Rx Mask requirements. For the training sequence simplicity and difficulty to capture CA inputs prior to training the CA inputs, the training mode described here uses a minimum of external commands to enter, train, and exit the Command Bus Training.

4.2.5 WCK2CK Leveling

4.2.5.1 WCK2CK Leveling Mode (write-leveling called in LPDDR4)

To adjust CK-to-WCK relationship and guarantee WCK2CK-Sync. operation, the LPDDR5 SDRAM provides a WCK2CK Leveling feature to compensate CK-to-WCK timing skew affecting WCK2CK-Sync. operation. The SDRAM compares the phase of the rising edge of WCK and the rising edge of CK, then asynchronously feeds back to the memory controller for the WCK2CK phase detection result. After finishing WCK2CK Leveling, tWCK2CK which means CK-to-WCK relationship is determined and WCK2CK-Sync. operation will be performed with the optimized margin.

4.2.9 WCK-DQ Training

The LPDDR5 SDRAM uses an un-matched WCK-DQ path to enable high speed performance and save power in the SDRAM. As a result, WCK is required to be trained to arrive at the DQ latch center-aligned with the Data eye. The SDRAM DQ receiver is located at the DQ pad and has a shorter internal delay in the SDRAM than does the WCK signal. The SDRAM DQ receiver will latch the data present on the DQ bus when WCK reaches the latch, and training is accomplished by delaying the DQ signals relative to WCK such that the Data eye arrives at the receiver latch centered on the WCK transition.

Prior to normal operation, the LPDDR5 SDRAM is required to be initialized. The Power-Up, Initialization, and Power-off Procedure section provides detailed information covering device initialization.

Source: JEDEC Standard No. 209-5A at pp. 2, 46, 75 & 90.

153. Furthermore, the Accused '835 Products generate command signals to access an integrated circuit component. For example, the JEDEC LPDDR5 specification, JESD209-5A, specifies CA (command signals), DQ (data signals), and WCK (sampling signals) are used to

access a standard LPDDR5 memory, the CA signals providing the command and address input according to the Command Truth Table, as illustrated below:

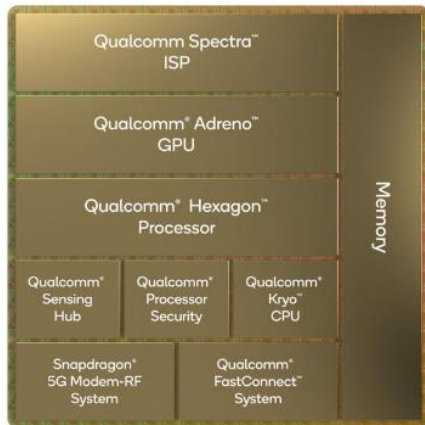
2.2.1 Pad Definition and Description

Table 1 — Pad Definition and Description

Symbol	Type	Description	Note
CK_t, CK_c	Input	Clock: CK_t and CK_c are differential clock inputs. All Double Data Rate (DDR) Command/Address inputs are sampled on both crossing points of CK_t and CK_c. The first crossing point is the rising(falling) edge of CK_t (CK_c) and second crossing point is falling(rising) edge of CK_t (CK_c). Single Data Rate (SDR) inputs, CS is sampled on the crossing point that is the rising(falling) edge of CK_t (CK_c).	
CS	Input	Chip Select: CS is part of the command code and is sampled on the rising(falling) edge of CK_t (CK_c) unless the device is in power-down or Deep Sleep mode where it becomes an asynchronous signal.	
CA[6:0]	Input	Command/Address Inputs: CA signals provide the Command and Address input according to the Command Truth Table	
DQ[15:0]	I/O	Data Input/Output: Bi-direction data bus.	
WCK[1:0]_t WCK[1:0]_c	Input	Data Clocks: WCK_t and WCK_c are differential clocks used for WRITE data capture and READ data output.	

SDRAM COMMAND
DESELECT (DES)
NO OPERATION (NOP)
POWER DOWN ENTRY (PDE)
ACTIVATE-1 (ACT-1)
ACTIVATE-2 (ACT-2)
PRECHARGE (PRE) (Per Bank, All Banks)
REFRESH (REF) (Per Bank, All Banks)
MASK WRITE (MWR)
WRITE (WR16 or WR)
WRITE32 (WR32)
READ (RD16 or RD)
READ32 (RD32)

Source: JEDEC Standard No. 209-5A at pp. 3, 165.



Source: <https://www.forbes.com/sites/marcochiappetta/2021/11/30/snapdragon-8-gen-1-the-qualcomm-mobile-platform-that-will-power-next-gen-android-flagship-phones/?sh=1dc07a4f37cd>

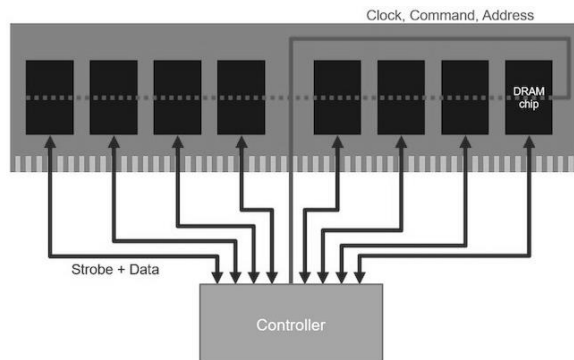
154. The Accused '835 Products access data signals to convey data for the integrated circuit component. For example, the Accused '835 Products, including the Motorola Edge+ smartphone (with the Qualcomm Snapdragon 8 Gen 1 Mobile Platform), support LPDDR5 memory and access data I/O via a bi-directional data bus to convey data for the LPDDR5 memory, as seen below:

2.2.1 Pad Definition and Description

Table 1 — Pad Definition and Description

Symbol	Type	Description	Note
CK_t, CK_c	Input	Clock: CK_t and CK_c are differential clock inputs. All Double Data Rate (DDR) Command/Address inputs are sampled on both crossing points of CK_t and CK_c. The first crossing point is the rising(falling) edge of CK_t (CK_c) and second crossing point is falling(rising) edge of CK_t (CK_c). Single Data Rate (SDR) inputs, CS is sampled on the crossing point that is the rising(falling) edge of CK_t (CK_c).	
CS	Input	Chip Select: CS is part of the command code and is sampled on the rising(falling) edge of CK_t (CK_c) unless the device is in power-down or Deep Sleep mode where it becomes an asynchronous signal.	
CA[6:0]	Input	Command/Address Inputs: CA signals provide the Command and Address input according to the Command Truth Table	
DQ[15:0]	I/O	Data Input/Output: Bi-direction data bus.	
WCK[1:0]_t WCK[1:0]_c	Input	Data Clocks: WCK_t and WCK_c are differential clocks used for WRITE data capture and READ data output.	

Source: JEDEC Standard No. 209-5A at 3.



Source: <https://www.signalintegrityjournal.com/blogs/8-for-good-measure/post/473-ddr-memory-interface-basics>

4.2.9 WCK-DQ Training

The LPDDR5 SDRAM uses an un-matched WCK-DQ path to enable high speed performance and save power in the SDRAM. As a result, WCK is required to be trained to arrive at the DQ latch center-aligned with the Data eye. The SDRAM DQ receiver is located at the DQ pad and has a shorter internal delay in the SDRAM than does the WCK signal. The SDRAM DQ receiver will latch the data present on the DQ bus when WCK reaches the latch, and training is accomplished by delaying the DQ signals relative to WCK such that the Data eye arrives at the receiver latch centered on the WCK transition.

Source: JEDEC Standard No. 209-5A at p. 90

155. The Accused '835 Products access sampling signals to control sampling of the data signals. For example, the Accused '835 Products, including the Motorola Edge+ smartphone (with the Qualcomm Snapdragon 8 Gen 1 Mobile Platform), support LPDDR5 memory and utilize WCK2CK Leveling, which accesses sampling signals to control sampling of the data signals, as illustrated below:

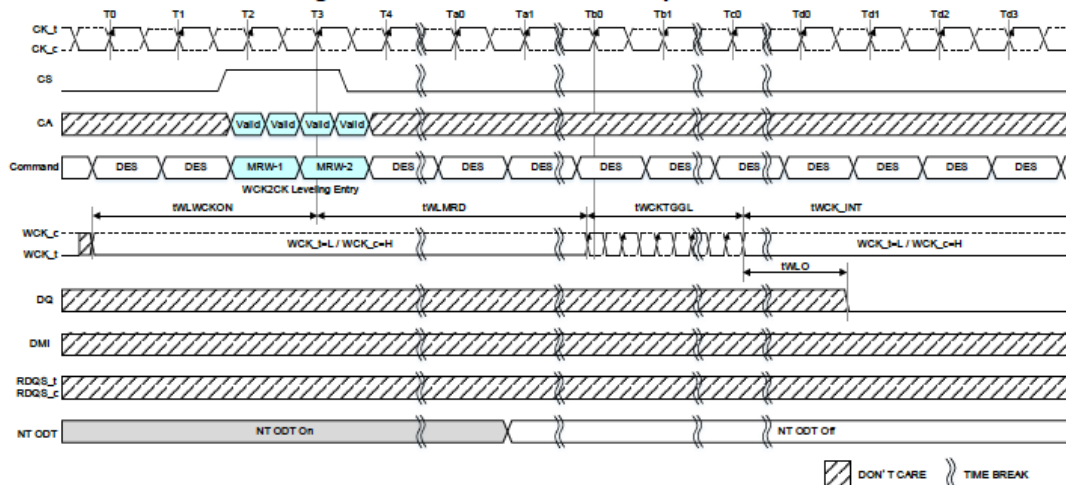
2.2.1 Pad Definition and Description

Table 1 — Pad Definition and Description

Symbol	Type	Description	Note
CK_t, CK_c	Input	Clock: CK_t and CK_c are differential clock inputs. All Double Data Rate (DDR) Command/Address inputs are sampled on both crossing points of CK_t and CK_c. The first crossing point is the rising(falling) edge of CK_t (CK_c) and second crossing point is falling(rising) edge of CK_t (CK_c). Single Data Rate (SDR) inputs, CS is sampled on the crossing point that is the rising(falling) edge of CK_t (CK_c).	
CS	Input	Chip Select: CS is part of the command code and is sampled on the rising(falling) edge of CK_t (CK_c) unless the device is in power-down or Deep Sleep mode where it becomes an asynchronous signal.	
CA[6:0]	Input	Command/Address Inputs: CA signals provide the Command and Address input according to the Command Truth Table	
DQ[15:0]	I/O	Data Input/Output: Bi-direction data bus.	
WCK[1:0]_t WCK[1:0]_c	Input	Data Clocks: WCK_t and WCK_c are differential clocks used for WRITE data capture and READ data output.	

Source: JEDEC Standard No. 209-5A at p. 3.

4.2.5.2 WCK2CK Leveling Procedure and Related AC parameters



Source: JEDEC Standard No. 209-5A at p. 76.

4.2.5 WCK2CK Leveling

4.2.5.1 WCK2CK Leveling Mode (write-leveling called in LPDDR4)

To adjust CK-to-WCK relationship and guarantee WCK2CK-Sync. operation, the LPDDR5 SDRAM provides a WCK2CK Leveling feature to compensate CK-to-WCK timing skew affecting WCK2CK-Sync. operation. The SDRAM compares the phase of the rising edge of WCK and the rising edge of CK, then asynchronously feeds back to the memory controller for the WCK2CK phase detection result. After finishing WCK2CK Leveling, tWCK2CK which means CK-to-WCK relationship is determined and WCK2CK-Sync. operation will be performed with the optimized margin.

Source: JEDEC Standard No. 209-5A at p. 75.

156. The Accused '835 Products systematically alter a phase shift of the command signals. For example, the Accused '835 Products, including the Motorola Edge+ smartphone (with the Qualcomm Snapdragon 8 Gen 1 Mobile Platform), follow the JEDEC 209-5A standard, which explains that the Command Bus training (CA training) requires phase adjustments of the CS and CA signals to be implemented at initialization (i.e., powerup), prior to normal operation, as seen below:

4.2.2 Command Bus Training

The LPDDR5 SDRAM command bus must be trained before enabling termination for high-frequency or mid-frequency operation. LPDDR5 SDRAM provides an internal VREF(CA) that default level is suitable for un-terminated, low-frequency operation, however the VREF(CA) must be trained to achieve suitable receiver voltage margin for terminated, high-frequency or mid-frequency operation. The training mode described here centers the internal VREF(CA) in the CA data eye and at the same time allows for timing adjustments of the CS and CA signals to meet their Rx Mask requirements. For the training sequence simplicity and difficulty to capture CA inputs prior to training the CA inputs, the training mode described here uses a minimum of external commands to enter, train, and exit the Command Bus Training.

Once LPDDR5 SDRAM has entered CBT mode by MRW, only DES and MRW for exiting CBT mode are allowed.

The LPDDR5 SDRAM supports two Command Bus Training modes and their feature is as follows. CBT mode is selected by MR13 OP[6] (CBT mode1: MR13 OP[6] = 0_B, CBT mode2: MR13 OP[6] = 1_B)

In multi-rank/channel system sharing the CA bus, the terminated die should be trained first, followed by the nonterminated die(s). See 7.6.4 for more information. For the WCK ODT setting in multi-rank/channel system, only one of SDRAM connected to a common WCK signal can set to CBT mode, the MR of Non-CBT trained SDRAM(s) is required to be set MR18 OP[2:0]=000_B.

The Corresponding DQ pins in this definition may differ depending on the package configuration. For example, in case of a package which contains Byte-mode devices, DQ[15:8] and DMI[1] balls can be connected to DQ[7:0] and DMI[0] pads of byte-mode device.

Source: JEDEC Standard No. 209-5A at p. 46.

Prior to normal operation, the LPDDR5 SDRAM is required to be initialized. The Power-Up, Initialization, and Power-off Procedure section provides detailed information covering device initialization.

Source: JEDEC Standard No. 209-5A at p. 2.

4.1.1 Voltage Ramp and Device Initialization (Cont'd)

- 7) Since LPDDR5 initial ZQ calibration is done automatically after ramp up, ZQ Latch command should be issued. After t_{ZQLAT} is satisfied (Th) the command bus (internal VREF(CA), CS, and CA) should be trained for high-speed operation by issuing MRW command (Command Bus Training Mode). This command is used to calibrate the SDRAM's internal VREF and align CS/CA with CK for high-speed operation. The LPDDR5 SDRAM will power-up with receivers configured for low-speed operations, and VREF(CA) set to a default factory setting. Normal SDRAM operation at clock speeds higher than t_{CKb} may not be possible until command bus training has been completed.

Source: JEDEC Standard No. 209-5A at p. 28.

157. The Accused '835 Products systematically alter a phase shift of the data signals. For example, the Accused '835 Products, including the Motorola Edge+ smartphone (with the Qualcomm Snapdragon 8 Gen 1 Mobile Platform), follow the JEDEC 209-5A standard, which explains that WCK-DQ training adjusts the timing (phase) of the data bus (data signals) during initialization (i.e., powerup), prior to normal operation, as illustrated below:

4.2.9 WCK-DQ Training

The LPDDR5 SDRAM uses an un-matched WCK-DQ path to enable high speed performance and save power in the SDRAM. As a result, WCK is required to be trained to arrive at the DQ latch center-aligned with the Data eye. The SDRAM DQ receiver is located at the DQ pad and has a shorter internal delay in the SDRAM than does the WCK signal. The SDRAM DQ receiver will latch the data present on the DQ bus when WCK reaches the latch, and training is accomplished by delaying the DQ signals relative to WCK such that the Data eye arrives at the receiver latch centered on the WCK transition.

The LPDDR5 SDRAM provides a Command-based FIFO Write/Read training operation using user specific pattern. Basically, DMI will be treated the same as DQs. It means that the Write Data send to FIFO for DMI by WFF command and these data can be read-out from FIFO for DMI by RFF command. On the other hand, DMI behavior is not same as DQs in some cases. Refer to the details about the special DMI behavior which are to be described later in this section.

Source: JEDEC Standard No. 209-5A at p. 90.

Prior to normal operation, the LPDDR5 SDRAM is required to be initialized. The Power-Up, Initialization, and Power-off Procedure section provides detailed information covering device initialization.

Source: JEDEC Standard No. 209-5A at p. 2.

4.1.1 Voltage Ramp and Device Initialization (Cont'd)

- 7) Since LPDDR5 initial ZQ calibration is done automatically after ramp up, ZQ Latch command should be issued. After t_{ZQLAT} is satisfied (Th) the command bus (internal VREF(CA), CS, and CA) should be trained for high-speed operation by issuing MRW command (Command Bus Training Mode). This command is used to calibrate the SDRAM's internal VREF and align CS/CA with CK for high-speed operation. The LPDDR5 SDRAM will power-up with receivers configured for low-speed operations, and VREF(CA) set to a default factory setting. Normal SDRAM operation at clock speeds higher than t_{CKb} may not be possible until command bus training has been completed.

Source: JEDEC Standard No. 209-5A at p. 28.

158. The Accused '835 Products systematically alter a phase shift of the sampling signals. For example, the Accused '835 Products, including the Motorola Edge+ smartphone (with

the Qualcomm Snapdragon 8 Gen 1 Mobile Platform), follow the JEDEC 209-5A standard, which explains that WCK2CK Leveling adjusts the timing (phase) of the WCK (sampling signals) during device initialization, prior to normal operation, as illustrated below:

4.2.5.1 WCK2CK Leveling Mode (write-leveling called in LPDDR4)

To adjust CK-to-WCK relationship and guarantee WCK2CK-Sync. operation, the LPDDR5 SDRAM provides a WCK2CK Leveling feature to compensate CK-to-WCK timing skew affecting WCK2CK-Sync. operation. The SDRAM compares the phase of the rising edge of WCK and the rising edge of CK, then asynchronously feeds back to the memory controller for the WCK2CK phase detection result. After finishing WCK2CK Leveling, tWCK2CK which means CK-to-WCK relationship is determined and WCK2CK-Sync. operation will be performed with the optimized margin.

The memory controller references WCK2CK-Sync feedback to adjust CK-to-WCK relationship for each WCK_t/WCK_c signal pair. All data bits (DQ[7:0] for WCK_t[0]/WCK_c[0], and DQ[15:8] for WCK_t[1]/WCK_c[1]) carry the training feedback to the controller.

CKR mode is required to be set 2:1 prior to entering WCK2CK leveling.

The LPDDR5 SDRAM enters into WCK2CK leveling mode when mode register MR18-OP[6] is set HIGH. When WCK2CK Leveling mode is entered, the state of the DQ pins is undefined. During WCK2CK Leveling mode, only DESELECT commands are allowed, or MRW command to exit the WCK2CK Leveling operation. Upon completion of the WCK2CK Leveling, the SDRAM exits from WCK2CK Leveling mode when MR18-OP[6] is reset LOW.

WCK2CK Leveling should be performed before write training.

WCK2CK Leveling examples are shown in Figure 42 and Figure 43, and the specific descriptions for the figures will be provided in the following section.

Source: JEDEC Standard No. 209-5A at p. 75.

Prior to normal operation, the LPDDR5 SDRAM is required to be initialized. The Power-Up, Initialization, and Power-off Procedure section provides detailed information covering device initialization.

Source: JEDEC Standard No. 209-5A at p. 2.

4.1.1 Voltage Ramp and Device Initialization (Cont'd)

- 7) Since LPDDR5 initial ZQ calibration is done automatically after ramp up, ZQ Latch command should be issued. After tZQLAT is satisfied (Th) the command bus (internal VREF(CA), CS, and CA) should be trained for high-speed operation by issuing MRW command (Command Bus Training Mode). This command is used to calibrate the SDRAM's internal VREF and align CS/CA with CK for high-speed operation. The LPDDR5 SDRAM will power-up with receivers configured for low-speed operations, and VREF(CA) set to a default factory setting. Normal SDRAM operation at clock speeds higher than tCKb may not be possible until command bus training has been completed.

Source: JEDEC Standard No. 209-5A at p. 28.

159. The Accused '835 Products systematically alter the phase shift of the command signals, the phase shift of the data signals, and the phase shift of the sampling signals to determine a valid operation range of the integrated circuit device, wherein the valid operation range includes an optimal operation point for the integrated circuit device. For example, the Accused '835 Products, including the Motorola Edge+ smartphone (with the Qualcomm Snapdragon 8 Gen 1

Mobile Platform), uses LPDDR5 memory operating at a maximum speed of 3200 MHz and requires LPDDR5 signal training to meet the tighter timing requirements of the LPDDR5 standard, notwithstanding delays in the memory controller and memory components introduced by myriad factors such as variations inherent in manufacturing processes, and varying conditions such as temperature and voltage. The training flow includes command bus training (including altering a phase shift of the command signals), WCK2CK Leveling (including altering a phase shift of the sampling signals) and WCK-DQ training (including altering a phase shift of the data signals).

4.2.2 Command Bus Training

- 4) At time t_{CAENT} later, LPDDR5 SDRAM can accept to input CA training pattern via CA bus.
- 5) To verify that the receiver has the correct VREF(CA) setting and to further train the CA eye relative to clock(CK), values latched at the receiver on the CA bus are asynchronously output to the DQ bus.
- 6) To exit Command Bus Training mode, drive DQ[7] LOW and after time $t_{DQ7LWCK} + t_{VREFCA_LONG}$ issue the MRW command to set MR16 OP[5:4] = 00_B. After time t_{MRD} the LPDDR5 SDRAM is ready for normal operation. After training exit, the LPDDR5 SDRAM will automatically switch back to the FSP-OP registers that were in use prior to training.

4.2.5.2 WCK2CK Leveling Procedure

4. Toggle WCK signal 7.5 cycles for WCK2CK phase detection. SDRAM may or may not capture the first rising edge of WCK_t due to an unstable first rising edge. Hence providing exactly 7.5 cycles of WCK signal input is required in every WCK input signal during WCK2CK training mode. SDRAM provides asynchronous feedback of last captured WCK2CK phase information during WCK toggles, on all the DQ bits after time t_{WLO} . DQ output is low if WCK phase is earlier than CK phase and high if WCK phase is later than CK phase. The controller must sample the phase relation result on DQ after satisfying t_{WLO} .
5. The feedback provided by the SDRAM is referenced by the controller to increment or decrement the WCK_t and WCK_c delay setting. The controller can adjust the WCK delay setting only when it drives WCK_t LOW and WCK_c HIGH to prevent any glitches in WCK signal. WCK search range from controller is defined as $t_{WCK2CK_leveling}$ ac parameter. Refer to the $t_{WCK2CK_leveling}$ value in Table 30.
6. Repeat step 4 through step 5 until the proper WCK_t/WCK_c delay is established.

4.2.9 WCK-DQ Training

The LPDDR5 SDRAM uses an un-matched WCK-DQ path to enable high speed performance and save power in the SDRAM. As a result, WCK is required to be trained to arrive at the DQ latch center-aligned with the Data eye. The SDRAM DQ receiver is located at the DQ pad and has a shorter internal delay in the SDRAM than does the WCK signal. The SDRAM DQ receiver will latch the data present on the DQ bus when WCK reaches the latch, and training is accomplished by delaying the DQ signals relative to WCK such that the Data eye arrives at the receiver latch centered on the WCK transition.

Source: JEDEC Standard No. 209-5A at 46-47, at pp. 77 & 90.

Prior to normal operation, the LPDDR5 SDRAM is required to be initialized. The Power-Up, Initialization, and Power-off Procedure section provides detailed information covering device initialization.

4.1.1 Voltage Ramp and Device Initialization

The following sequence shall be used to power up the LPDDR5 SDRAM. Unless specified otherwise, these steps are mandatory.

4.1.1 Voltage Ramp and Device Initialization (Cont'd)

- 7) Since LPDDR5 initial ZQ calibration is done automatically after ramp up, ZQ Latch command should be issued. After tZQLAT is satisfied (Th) the command bus (internal VREF(CA), CS, and CA) should be trained for high-speed operation by issuing MRW command (Command Bus Training Mode). This command is used to calibrate the SDRAM's internal VREF and align CS/CA with CK for high-speed operation. The LPDDR5 SDRAM will power-up with receivers configured for low-speed operations, and VREF(CA) set to a default factory setting. Normal SDRAM operation at clock speeds higher than tCKb may not be possible until command bus training has been completed.

NOTE 1 The command bus training MRW command uses the CA bus as inputs for the calibration data stream, and outputs the results asynchronously on the DQ bus. See 4.2.2 for information on how to enter/exit the training mode.

- 8) After command bus training, DRAM controller must perform WCK2CK leveling. WCK2CK leveling mode is enabled when MR18-OP[6] is high (Ti). See 4.2.5.2 for detailed description of WCK2CK leveling entry and exit sequence. After finishing WCK2CK Leveling, tWCK2CK which means CK-to-WCK relationship is determined and WCK2CK-Sync. operation will be performed with the optimized margin.
- 9) After WCK2CK leveling, the DQ Bus (internal VREF(DQ), WCK, and DQ) should be trained for high-speed operation using the training commands (RD FIFO / WT FIFO / RD DQ Calibration) described in command truth table and by issuing MRW commands to adjust VREF(DQ)(Ti). The LPDDR5 SDRAM will power-up with receivers configured for low-speed operations and VREF(DQ) set to a default factory setting. Normal SDRAM operation at clock speeds higher than tCKb should not be attempted until DQ Bus training has been completed. The Read DQ Calibration command is used together with FIFO Write/Read commands to train DQ bus without disturbing the memory array contents. See 4.2.9 for detailed DQ Bus Training sequence.

Source: JEDEC Standard No. 209-5A at 46-47, at pp. 2, 28.

160. Additionally, Defendant has been, and currently is, an active inducer of infringement of the '835 patent under 35 U.S.C. § 271(b) and a contributory infringer of the '835 patent under 35 U.S.C. § 271(c).

161. Defendant has actively induced, and continues to actively induce, infringement of the '835 patent by causing others to use, offer for sale, or sell in the United States, products or services covered by the '835 patent, including but not limited to the '835 Accused Products and any other products or services that include Qualcomm processors and LPDDR4, LPDDR4X, or LPDDR5 memory, or products and services with the same or substantially similar functionality to that described above. Defendant provides these products and services to others, such as customers, resellers, partners, and end-users, who, in turn, use, provision for use, offer for sale, or sell those products and services, which directly infringe the '835 patent as described above. Defendant's inducement includes requiring memory chips within the Accused Products to be compliant with the JESD209-5A JEDEC standard for Low Power Double Data Rate 5 (LPDDR5), in which the

calibration of timing relationships between command signals, data signals and sampling signals described above is mandatory, and advertising and promoting such compliance to its customers, partners, re-sellers and the like, including the promotion, directions and instructions found at one or more of the following links, the provision of which has been on-going since the filing of the First Amended Complaint in case 6:23-cv-0068-ADA and the content of which is specifically illustrated above:

- <https://www.motorola.com/us/smartphones-motorola-edge-plus-gen-2/p?skuId=774>
- [https://www.lenovo.com/us/en/p/laptops/thinkpad/thinkpadx/thinkpad-x13s-\(13-inch-snapdragon\)/21bx0008us](https://www.lenovo.com/us/en/p/laptops/thinkpad/thinkpadx/thinkpad-x13s-(13-inch-snapdragon)/21bx0008us)
- [https://www.lenovo.com/us/en/p/laptops/lenovo/windows-edu-laptops/lenovo-10w-\(10-inch-qlc\)/82st0002us](https://www.lenovo.com/us/en/p/laptops/lenovo/windows-edu-laptops/lenovo-10w-(10-inch-qlc)/82st0002us)
- <https://investors.micron.com/news-releases/news-release-details/microns-low-power-ddr5-dram-boosts-performance-and-consumer>
- <https://www.micron.com/about/blog/2021/june/lpddr5-brings-flagship-features-to-a-wider-range-of-phones>

162. Defendant has contributed to, and continues to contribute to, the infringement of the '835 patent by others by knowingly providing one or more components, for example the Qualcomm processors and LPDDR4, LPDDR4X and LPDDR5 complaint memory included in the Accused Products, a portion thereof, and/or the software/hardware modules responsible for the accused functionality described herein, that, when installed, configured, and used result in systems that, as intended by Defendant described above, directly infringe one or more claims of the '835 patent.

163. Defendant knew of the '835 patent, or should have known of the '835 patent, but was willfully blind to its existence. Upon information and belief, Defendant had actual knowledge of the '835 patent at least as early as the filing of the First Amended Complaint in case 6:23-cv-0068-ADA, or alternatively, at least as early as Defendant's receipt of this Complaint.

Alternatively, upon information and belief, Defendant has had knowledge of the '835 patent since the service upon Defendant of the Complaint in this action.

164. By the time of trial, Defendant will or should have known and intended (since receiving such notice) that its continued actions would infringe and would actively induce and contribute to the infringement of the '835 patent.

165. Defendant has committed, and continues to commit, contributory infringement by selling products and services that directly infringe the '835 patent when used by a third party, such as the Accused '835 Products, and that are a material part of the invention, knowing them to be especially made or adapted for use in infringement of the '835 patent and not staple articles or commodities of commerce suitable for substantial non-infringing use.

166. As a result of Defendant's acts of infringement, IV has suffered and will continue to suffer damages in an amount to be determined at trial.

PRAYER FOR RELIEF

IV requests that the Court enter judgment as follows:

- (A) that Defendant has infringed the '140 patent;
- (B) that Defendant has infringed the '016 patent;
- (C) that Defendant has infringed the '443 patent;
- (D) that Defendant has infringed the '439 patent;
- (E) that Defendant has infringed the '835 patent;
- (F) awarding damages sufficient to compensate IV for Defendant's infringement under 35 U.S.C. § 284;
- (G) finding this case exceptional under 35 U.S.C. § 285 and awarding IV its reasonable attorneys' fees;

- (H) awarding IV its costs and expenses incurred in this action;
- (I) awarding IV prejudgment and post-judgment interest; and
- (J) granting IV such further relief as the Court deems just and appropriate.

DEMAND FOR JURY TRIAL

IV demands trial by jury of all claims so triable under Federal Rule of Civil Procedure 38.

Dated: April 26, 2023.

Respectfully submitted,

/s/Karl Rupp

State Bar No. 24035243

SOREY & HOOVER, LLP

100 N. 6TH Street, Ste. 502

Waco, Texas 76701

Tel: (903) 230-5600

Fax: (903) 230-5656

krupp@soreylaw.com

Paul J. Hayes

phayes@princelobel.com

Matthew D. Vella

mvella@princelobel.com

Robert R. Gilman

rgilman@princelobel.com

Jonathan DeBlois

jdeblois@princelobel.com

Brian Seeve

bseeve@princelobel.com

PRINCE LOBEL TYE LLP

One International Place, Suite 3700

Boston, MA 02110

Tel: (617) 456-8000

COUNSEL FOR PLAINTIFFS