

# Scale in the Mega Data Center: Market Shift to Non-Blocking Network Inside Data Center Building

## Table of Contents

### Mega Data Center Scale: Executive Summary

Scale is a vital part of the technology used to support next generation data centers. The study is targeted to C-level executives that need to move quickly and surely to improve IT. Automation of IT depends on understanding the business market opportunity from an independent perspective. Vendors are smart but they are committed to the technology they are pushing, the Sea Change Series from WinterGreen Research is able to provide a perspective not available anywhere else because of the comprehensive in-depth, yet simple look at the technology impact on business.

Extreme scale is what brings enough pathways inside a Mega data center to create a non-blocking (CLOS) networked server architecture. Non-blocking network architecture benefits the business because it permits launching thousands of virtual servers on demand at the application layer. In this manner, innovation can be made to happen quickly.

### Sea Change Series: Scale in the Mega Data Center

The data center is part of sea change because the cabling and servers are too slow to handle the mountains of data coming from a shift to everything digital. Those companies that are able to adapt can benefit from the ability to implement innovative projects quickly and inexpensively, stimulating rapid growth and enormous competitive advantage.

Amazon, Google, Microsoft, Facebook	2
Aim to Realign IT Cost Structure	3
Scale Matters	4
Table of Contents	5

Table of Contents and List of Figures

Facebook Mega Datacenter Physical Infrastructure	13
Facebook Automation of Mega Data Center Process	14
Facebook Altoona Data Center Networking Fabric	15
Facebook Altoona Cloud Mega Data Center	16
Facebook Altoona Data Center Innovative Networking Fabric Depends on Scale	17
Facebook Fabric Operates Inside the Data Center	18
Facebook Fabric	19
Exchange Of Data Between Servers Represents A Complex Automation Of Process	20
Applications Customized For Each User	21
Machine-To-Machine Management of Traffic Growth	22
Facebook Data Center Fabric Network Topology	23
Building-Wide Connectivity	24
Highly Modular Design Allows Users To Quickly Scale Capacity In Any Dimension	25
Back-End Service Tiers And Applications	26
Scaling Up As a Basic Function Of The Mega Data Center Network	27
Facebook Fabric Next-Generation Data Center Network Design: Pod Unit of Network	28
Mega Data Center Server Pods	29
Facebook Sample Pod: Unit of Network	30
Non-Blocking Network Architecture	31
Data Center Auto Discovery	36
Facebook Large-Scale Network	37
Rapid Deployment Architecture	38
Facebook Expedites Provisioning And Changes	39
Google Douglas County Mega Data Center	40
Google Data Center Efficiency Measurements	41
Google Programmable Access To Network Stack	42
Google Software Defined Networking (SDN)-Supports Scale and Automation	43
Google Compute Engine Load Balancing	44
Google Compute Engine Load Balanced Requests Architecture	45
Google Compute Engine Load Balancing Scaling	46
Google Switches Provide Scale-Out: Server and Storage Expansion	47
Google Uses Switches and Routers Deployed in Fabrics	48
Google Mega Data Center Multi-pathing	49
Google Mega Data Center Multi-pathing: Routing Destinations	50
Google Clos Topology Network Capacity Scalability	51
Google Aggregation Switches Are Lashed Together Through a Set Of Non-Blocking Spine Switches	52
Google Network Called Jupiter	53

Table of Contents and List of Figures

Microsoft Cloud Data Center Multi-Tenant Containers	54
Microsoft Azure Running Docker Containers	55
Microsoft Data Center, Dublin, 550,000 Sf	56
Microsoft Builds Intelligent Cloud Platform	57
Microsoft Crafts Homegrown Linux For Azure Switches	59
Microsoft Azure Has Scale	60
Microsoft Azure Stack Hardware Foundation	62
Microsoft Azure Stack Key Systems Partners: Cisco Systems, Lenovo, Fujitsu, and NEC	63
Microsoft Gradual Transformation From A Platform Cloud To A Broader Offering	
Leveraging Economies of Scale	64
Microsoft Contributing to Open Systems	65
Microsoft Mega Data Center Supply Chain	66
Microsoft Leverages Open Compute Project to Bring Benefit to Enterprise Customers	67
Microsoft Assists Open Compute to Close The Loop On The Hardware Side	68
Microsoft Project Olympus Modular And Flexible	69
Microsoft Azure 70	
Microsoft Azure Active Directory Has Synchronization	71
Microsoft Azure Has Scale	72
Mega Data Center Different from the Hyperscale Cloud	73
Mega Data Center Scaling	74
Mega Data Center Automatic Rules and Push-Button Actions	75
Amazon Capex for Cloud 2.0 Mega Data Centers	76
AWS Server Scale	77
Amazon North America	78
Innovation a Core Effort for Amazon	80
Amazon Offers the Richest Services Set	81
AWS Server Scale	81
On AWS, Customers Architect Their Applications	82
AWS Scale to Address Network Bottleneck	83
Networking A Concern for AWS Solved by Scale	84
AWS Regions and Network Scale	85
AWS Datacenter Bandwidth	88
Amazon (AWS) Regional Data Center	89
Map of Amazon Web Service Global Infrastructure	90
Rows of Servers Inside an Amazon (AWS) Data Center	91
Amazon Capex for Mega Data Centers	92

Table of Contents and List of Figures

Amazon Addresses Enterprise Cloud Market, Partnering With VMware Making Individual Circuits And Devices Unimportant Is A Primary Aim Of Fabric Architecture	92
Google Clos Network Architecture Topology Allows the Building a Non-Blocking Network Using Small Switches	93
You Have To Hit A Certain Scale Before Clos Networks Work Clos Network	94
Digital Data Expanding Exponentially, Global IP Traffic Passes Zettabyte (1000 Exabytes) Threshold	95
Summary: Economies of Scale	96
<b>WinterGreen Research,</b>	<b>101</b>
WinterGreen Research Methodology	102

## List of Figures

Enterprise Data Center as a Bottleneck: Scale Supports Non Blocking Network Inside Building and More Efficient Processing	
Figure 1. Slow Growth Companies Do Not Have Data Center Scale	2
Figure 2. Mega Data Center Fabric Implementation	3
Figure 3. Facebook Schematic Fabric-Optimized Datacenter Physical Topology	13
Figure 4. Facebook Automation of Mega Data Center Process	14
Figure 5. Facebook Altoona Positioning Of Global Infrastructure	15
Figure 6. FaceBook Equal Performance Paths Between Servers	16
Figure 7. FaceBook Data Center Fabric Depends on Scale	17
Figure 8. Facebook Fabric Operates Inside the Data Center, Fabric Is The Whole Data Center	18
Figure 9. Fabric Switches and Top of Rack Switches, Facebook Took a Disaggregated Approach	19

Table of Contents and List of Figures

Figure 10. Exchange Of Data Between Servers Represents A Complex Automation Of Process	20
Figure 11. Samsung Galaxy J3	21
Figure 12. Facebook Back-End Service Tiers And Applications Account for Machine-To-Machine Traffic Growth	22
Figure 1. Facebook Data Center Fabric Network Topology	23
Figure 13. Implementing building-wide connectivity	24
Figure 14. Modular Design Allows Users To Quickly Scale Capacity In Any Dimension	25
Figure 15. Facebook Back-End Service Tiers And Applications Functions	26
Figure 16. Using Fabric to Scale Capacity	27
Figure 17. Facebook Fabric: Pod Unit of Network	28
Figure 18. Server Pods Permit An Architecture Able To Implement Uniform High-Performance Connectivity	29
Figure 19. Non-Blocking Network Architecture	31
Figure 20. Facebook Automation of Cloud 2.0 Mega Data Center Process	32
Figure 21. Facebook Creating a Modular Cloud 2.0 mega data center Solution	33
Figure 22. Facebook Cloud 2.0 mega data center Fabric High-Level Settings Components	34
Figure 23. Facebook Mega Data Center Fabric Unattended Mode	35
Figure 24. Facebook Data Center Auto Discovery Functions	36
Figure 25. Facebook Automated Process Rapid Deployment Architecture	38
Figure 26. Google Douglas County Cloud 2.0 Mega Data Center	40
Figure 27. Google Data Center Efficiency Measurements	41
Figure 28. Google Andromeda Cloud High-Level Architecture	42
Figure 29. Google Andromeda Software Defined Networking (SDN)-Based Substrate Functions	43
Figure 30. Google Compute Engine Load Balancing Functions	44

Table of Contents and List of Figures

Figure 31. Google Compute Engine Load Balanced Requests Architecture	45
Figure 32. Google Compute Engine Load Balancing Scaling	46
Figure 33. Google Traffic Generated by Data Center Servers	47
Figure 34. Google Mega Data Center Multi-pathing: Implementing Lots And Lots Of Paths Between Each Source And Destination	49
Figure 35. Google Mega Data Center Multi-pathing: Routing Destinations	50
Figure 36. Google Builds Own Network Switches And Software	50
Figure 37. Google Clos Topology Network Capacity Scalability	51
Figure 38. Schematic Fabric-Optimized Facebook Datacenter Physical Topology	52
Figure 39. Google Jupiter Network Delivers 1.3 Pb/Sec Of Aggregate Bisection Bandwidth Across A Datacenter	53
Figure 40. Microsoft Azure Cloud Software Stack Hyper-V hypervisor	54
Figure 41. Microsoft Azure Running Docker Containers	55
Figure 42. Microsoft Data Center, Dublin, 550,000 Sf	56
Figure 43. Microsoft-Azure-Stack-Block-Diagram	60
Figure 44. Microsoft Azure Stack Architecture	62
Figure 45. Microsoft Data Centers	66
Figure 46. Microsoft Open Hardware Design: Project Olympus	67
Figure 47. Microsoft Open Compute Closes That Loop On The Hardware Side	68
Figure 48. Microsoft Olympus Product	69
Figure 49. Microsoft Azure Has Scale	72
Figure 50. Mega Data Center Cloud vs. Hyperscale Cloud	73
Figure 51. Amazon Web Services	76
Figure 52. Amazon North America Map	78
Figure 53. Amazon North America List of Locations	79
Figure 54. Woods Hole Bottleneck: Google Addresses Network Bottleneck with Scale	83
Figure 55. Example of AWS Region	85

Table of Contents and List of Figures

Figure 56. Example of AWS Availability Zone	86
Figure 57. Example of AWS Data Center	87
Figure 58. AWS Network Latency and Variability	88
Figure 59. Amazon (AWS) Regional Data Center	89
Figure 60. A Map of Amazon Web Service Global Infrastructure	90
Figure 61. Rows of Servers Inside an Amazon (AWS) Data Center	91
Figure 62. Clos Network	96
Figure 63. Data Center Technology Shifting	97
Figure 64. Data Center Technology Shift	98