

Contents

Acknowledgements	i
Abstract (English/German/French)	iii
List of figures	xiii
List of tables	xv
1 Introduction	1
1.1 Thesis Statement	4
1.2 Thesis Contributions	4
1.2.1 VNTOR: Rack-Based Network Virtualization	5
1.2.2 CRISS-CROSS: A Hierarchical, Reconfigurable Data Center Topology	6
1.3 Thesis Roadmap	7
1.3.1 Bibliographic Notes	8
2 Background	9
2.1 Data Center Designs	9
2.1.1 Cloud Computing	9
2.1.2 Common Topologies	12
2.1.3 Alternative Topologies	15
2.1.4 Centralized Control	15
2.2 Switch Architecture	17
2.3 The Locality Principle	19
2.3.1 Traffic Patterns	22
3 VNTOR: Network Virtualization at the Top-of-Rack Switch	25
3.1 Introduction	25
3.2 Background	29
3.2.1 Security-Group Abstraction	29
3.2.2 OpenStack and OVS Enforcement	29
3.2.3 SR-IOV and Security Groups	30
3.2.4 Validation of the Status Quo	32
3.3 Our Proposal: Move to the ToR	33
3.4 System Design	35

3.4.1	Platform	35
3.4.2	Local vs. Remote Control Plane	35
3.4.3	Architecture	37
3.4.4	Cache Management	38
3.4.5	Software Forwarding	40
3.4.6	OpenStack Integration	41
3.5	Prototype Implementation	41
3.5.1	Hardware	41
3.5.2	Software	42
3.5.3	Traffic Separation	43
3.6	Evaluation	43
3.6.1	Experimental Setup	44
3.6.2	Baseline: Static Workload	45
3.6.3	Churn and Breaking Point	48
3.6.4	Trace-based Study	50
3.7	Related Work	54
3.8	Conclusion	55
4	CRISS-CROSS: Dissolving data-center pods while maintaining the hierarchy of a data-center fabric	57
4.1	Introduction	57
4.2	Background	61
4.2.1	Clos Topologies	61
4.2.2	Oversubscription and Scaling	62
4.2.3	Optical Circuit Switches	63
4.3	Problem statement	64
4.4	Design	65
4.4.1	CRISS-CROSS Topology	65
4.4.2	High-level Algorithm	68
4.4.3	Routing	68
4.4.4	Connectivity between ToRs	70
4.4.5	Deployability and Physical Layout	71
4.5	Implementation	71
4.5.1	Candidate Selection	71
4.5.2	Victim Selection	72
4.5.3	Link Change Procedure	72
4.5.4	Link-level Detection / Configuration	73
4.6	Evaluation	73
4.6.1	Simulation Environment	74
4.6.2	Experimental Methodology	74
4.6.3	Effect of Reconfiguration	76
4.6.4	Spine Utilization	81
4.6.5	Sustainable Load	83

4.6.6	Reconfiguration Rate and Flow Completion Times	83
4.6.7	Sensitivity to Topology Size	85
4.7	Related Work	91
4.8	Discussion	93
4.9	Conclusion	94
5	Discussion	95
5.1	Future work	95
5.2	Conclusion	97
	Bibliography	99
	Curriculum Vitae	113