



Contents

Foreword	xxi
Preface	xxv
Acknowledgments	xxx
About the Authors	xxxv
Part I Introduction	
Chapter 1 Introduction to DTrace	1
What Is DTrace?	1
Why Do You Need It?	1
Capabilities	2
Dynamic and Static Probes	4
DTrace Features	4
A First Look	6
Overview	8
Consumers	9
Probes	10
Providers	11
Predicates	13

Actions	13
Aggregations	13
D Language	14
Architecture	16
Summary	17
Chapter 2 D Language	19
D Language Components	20
Usage	20
Program Structure	21
Probe Format	21
Predicates	22
Actions	23
Probes	23
Wildcards	23
BEGIN and END	24
profile and tick	24
syscall Entry and Return	25
Variables	26
Types	26
Operators	27
Scalar	28
Associative Arrays	29
Structs and Pointers	29
Thread Local	30
Clause Local	30
Built-in	31
Macro	32
External	33
Aggregations	33
Types	34
quantize()	34
lquantize()	35

trunc() and clear()	36
normalize()	36
printa()	36
Actions	37
trace()	37
printf()	38
tracemem()	39
copyin()	39
stringof() and copyinstr()	39
strlen() and strjoin()	40
stack(), ustack(), and jstack()	40
sizeof()	41
exit()	41
Speculations	41
Translators	42
Others	42
Options	43
Example Programs	44
Hello World	44
Tracing Who Opened What	44
Tracing fork() and exec()	45
Counting System Calls by a Named Process	45
Showing Read Byte Distributions by Process	45
Profiling Process Names	46
Timing a System Call	47
Snoop Process Execution	48
Summary	49
Part II	Using DTrace
Chapter 3	System View
	51
	Start at the Beginning
	52
	System Methodology
	53
	System Tools
	54

Observing CPUs	56
CPU Strategy	56
CPUs and Interrupts	85
CPU Events	88
CPU Summary	94
Observing Memory	95
Memory Strategy	95
Memory Checklist	96
Memory Providers	96
Memory One-Liners	97
Memory Analysis	98
User Process Memory Activity	101
Kernel Memory	118
Memory Summary	124
Observing Disk and Network I/O	125
I/O Strategy	125
I/O Checklist	125
I/O Providers	126
I/O One-Liners	127
I/O Analysis	128
Disk I/O	134
Network I/O	141
Summary	148
Chapter 4 Disk I/O	151
Capabilities	152
Disk I/O Strategy	154
Checklist	155
Providers	156
io Provider	157
fbt Provider	163
One-Liners	165
One-Liner Examples	166

Scripts	172
io Provider Scripts	173
SCSI Scripts	211
SATA Scripts	236
IDE Scripts	250
SAS Scripts	259
Case Studies	269
Shouting in the Data Center: A Personal Case Study (Brendan)	269
DTracing an Unfamiliar I/O Driver (SATA)	273
Conclusion	290
Summary	290
Chapter 5 File Systems	291
Capabilities	292
Logical vs. Physical I/O	295
Strategy	295
Checklist	296
Providers	297
fsinfo Provider	298
io Provider	300
One-Liners	300
One-Liners: syscall Provider Examples	304
One-Liners: vminfo Provider Examples	308
One-Liners: fsinfo Provider Examples	308
One-Liners: sdt Provider Examples	312
Scripts	313
Syscall Provider	315
fsinfo Scripts	327
VFS Scripts	335
UFS Scripts	351
ZFS Scripts	357
HFS+ Scripts	370

PCFS Scripts	375
HSFS Scripts	376
UDFS Scripts	378
NFS Client Scripts	379
TMPFS Scripts	385
Case Study	387
ZFS 8KB Mirror Reads	387
Conclusion	397
Summary	397
Chapter 6 Network Lower-Level Protocols	399
Capabilities	400
Strategy	402
Checklist	403
Providers	404
mib Provider	405
ip Provider	408
Network Providers	411
fbt Provider	415
One-Liners	422
Scripts	445
Socket Scripts	447
IP Scripts	469
TCP Scripts	481
UDP Scripts	517
ICMP Scripts	521
XDR Scripts	529
Ethernet Scripts	533
Common Mistakes	548
Receive Context	548
Send Context	550
Packet Size	553
Stack Reuse	554
Summary	555

Chapter 7	Application-Level Protocols	557
	Capabilities	558
	Strategy	558
	Checklist	559
	Providers	560
	fbt Provider	561
	pid Provider	562
	One-Liners	563
	Scripts	574
	NFSv3 Scripts	576
	NFSv4 Scripts	592
	CIFS Scripts	599
	HTTP Scripts	609
	DNS Scripts	621
	FTP Scripts	625
	iSCSI Scripts	633
	Fibre Channel Scripts	646
	SSH Scripts	649
	NIS Scripts	663
	LDAP Scripts	664
	Multiscripts	666
	Summary	668
Chapter 8	Languages	669
	Capabilities	671
	Strategy	672
	Checklist	674
	Providers	675
	Languages	676
	Assembly	677
	C	679
	User-Land C	680
	Kernel C	681
	Probes and Arguments	681

Struct Types	682
Includes and the Preprocessor	683
C One-Liners	684
C One-Liners Selected Examples	687
See Also	688
C Scripts	689
C++	689
Function Names	690
Object Arguments	690
Java	691
Example Java Code	693
Java One-Liners	693
Java One-Liners Selected Examples	694
Java Scripts	696
See Also	705
JavaScript	705
Example JavaScript Code	707
JavaScript One-Liners	708
JavaScript One-Liners Selected Examples	709
JavaScript Scripts	712
See Also	718
Perl	719
Example Perl Code	720
Perl One-Liners	720
Perl One-Liners Selected Examples	721
Perl Scripts	722
PHP	731
Example PHP Code	733
PHP One-Liners	734
PHP One-Liners Selected Examples	735
PHP Scripts	736
Python	740
Example Python Code	741

Python One-Liners	741
Python One-Liners Selected Examples	742
Python Scripts	744
Ruby	751
Example Ruby Code	752
Ruby One-Liners	753
Ruby One-Liners Selected Examples	753
Ruby Scripts	755
See Also	762
Shell	764
Example Shell Code	765
Shell One-Liners	765
Shell One-Liners Selected Examples	766
Shell Scripts	768
See Also	774
Tcl	774
Example Tcl Code	776
Tcl One-Liners	776
Tcl One-Liners Selected Examples	777
Tcl Scripts	778
Summary	782
Chapter 9 Applications	783
Capabilities	784
Strategy	784
Checklist	786
Providers	787
pid Provider	788
cpc Provider	791
See Also	793
One-Liners	793
One-Liner Selected Examples	798

Scripts	804
procsnoop.d	804
procsystime	806
uoncpu.d	808
uoffcpu.d	809
plockstat	811
kill.d	813
sigdist.d	814
threaded.d	815
Case Studies	817
Firefox idle	817
Xvnc	824
Summary	832
Chapter 10 Databases	833
Capabilities	834
Strategy	835
Providers	836
MySQL	837
One-Liners	838
One-Liner Selected Examples	840
Scripts	841
See Also	850
PostgreSQL	851
One-Liners	853
One-Liner Selected Examples	854
Scripts	854
See Also	858
Oracle	858
Examples	858
Summary	865

Part III Additional User Topics

Chapter 11 Security	867
Privileges, Detection, and Debugging	867
DTrace Privileges	868
DTrace-Based Attacks	869
Sniffing	869
Security Audit Logs	870
HIDS	871
Policy Enforcement	871
Privilege Debugging	872
Reverse Engineering	874
Scripts	875
sshkeysnoop.d	875
shellsnoop	878
keylatency.d	882
cuckoo.d	884
watchexec.d	886
nosetuid.d	888
nosnoopforyou.d	890
networkwho.d	891
Summary	892
Chapter 12 Kernel	893
Capabilities	894
Strategy	896
Checklist	897
Providers	897
fbt Provider	898
Kernel Tracing	903
Kernel Memory Usage	908
Anonymous Tracing	917
One-Liners	918
One-Liner Selected Examples	925

Scripts	932
intrstat	932
lockstat	934
koncpu.d	937
koffcpu.d	938
taskq.d	939
priclass.d	941
cswstat.d	943
putnexts.d	944
Summary	945
Chapter 13 Tools	947
The DTraceToolkit	948
Locations	948
Versions	949
Installation	949
Scripts	949
Script Example: cpuwalk.d	957
Chime	962
Locations	962
Examples	963
DTrace GUI Plug-in for NetBeans and Sun Studio	966
Location	966
Examples	966
DLight, Oracle Solaris Studio 12.2	966
Locations	969
Examples	969
Mac OS X Instruments	971
Locations	972
Examples	972
Analytics	973
The Problem	973
Solving the Problem	974

Contents	xvii
Toward a Solution	975
Appliance Analytics	976
Summary	985
Chapter 14 Tips and Tricks	987
Tip 1: Known Workloads	987
Tip 2: Write Target Software	989
Tip 3: Use grep to Search for Probes	991
Tip 4: Frequency Count	991
Tip 5: Time Stamp Column, Postsort	992
Tip 6: Use Perl to Postprocess	993
Tip 7: Learn Syscalls	994
Tip 8: timestamp vs. vtimestamp	995
Tip 9: profile:::profile-997 and Profiling	996
Tip 10: Variable Scope and Use	997
Thread-Local Variables	997
Clause-Local Variables	998
Global and Aggregation Variables	999
Tip 11: strlen() and strcmp()	999
Tip 12: Check Assumptions	1000
Tip 13: Keep It Simple	1001
Tip 14: Consider Performance Impact	1001
Tip 15: drops and dynvardrops	1003
Tip 16: Tail-Call Optimization	1003
Further Reading	1003
Appendix A DTrace Tunable Variables	1005
Appendix B D Language Reference	1011
Appendix C Provider Arguments Reference	1025
Providers	1025
Arguments	1038
bufinfo_t	1038

devinfo_t	1038
fileinfo_t	1038
cpuinfo_t	1039
lwpsinfo_t	1039
psinfo_t	1039
conninfo_t	1040
pktinfo_t	1040
csinfo_t	1040
ipinfo_t	1040
ifinfo_t	1041
ipv4info_t	1041
ipv6info_t	1041
tcpinfo_t	1042
tcpsinfo_t	1042
tcplsinfo_t	1043
Appendix D DTrace on FreeBSD	1045
Enabling DTrace on FreeBSD 7.1 and 8.0	1045
DTrace for FreeBSD: John Birrell	1047
Appendix E USDT Example	1051
USDT Bourne Shell Provider	1052
Compared to SDT	1052
Defining the Provider	1052
Adding a USDT Probe to Source	1053
Stability	1055
Case Study: Implementing a Bourne Shell Provider	1057
Where to Place the Probes	1059
Appendix F DTrace Error Messages	1063
Privileges	1063
Message	1063
Meaning	1063
Suggestions	1064

Drops	1064
Message	1064
Meaning	1064
Suggestions	1064
Aggregation Drops	1065
Message	1065
Meaning	1065
Suggestions	1065
Dynamic Variable Drops	1066
Message	1066
Meaning	1066
Suggestions	1066
Invalid Address	1066
Message	1066
Meaning	1066
Suggestions	1067
Maximum Program Size	1067
Message	1067
Meaning	1067
Suggestions	1067
Not Enough Space	1068
Message	1068
Meaning	1068
Suggestions	1068
Appendix G DTrace Cheat Sheet	1069
Synopsis	1069
Finding Probes	1069
Finding Probe Arguments	1070
Probes	1070
Vars	1070
Actions	1071
Switches	1071

Pragmas	1071
One-Liners	1072
Bibliography	1073
Suggested Reading	1073
Vendor Manuals	1075
FreeBSD	1075
Mac OS X	1075
Solaris	1076
Glossary	1077
Index	1089