



Preview of Rules-based Scalable Vector Extension (SVE) Supplement

Document number	AES0047
Document quality	BETA
Document version	Aa
Document confidentiality	Non-confidential
Document build information	Printed on: February 19, 2021.

Copyright © 2021 Arm Limited or its affiliates. All rights reserved.

Draft

Release information

Date	Version	Changes
2021/Feb/19	Non-Confidential • Beta Beta A.a	

Draft

Non-Confidential Proprietary Notice

This document is protected by copyright and other related rights and the practice or implementation of the information contained in this document may be protected by one or more patents or pending patent applications. No part of this document may be reproduced in any form by any means without the express prior written permission of Arm. **No license, express or implied, by estoppel or otherwise to any intellectual property rights is granted by this document unless specifically stated.**

Your access to the information in this document is conditional upon your acceptance that you will not use or permit others to use the information for the purposes of determining whether implementations infringe any third party patents.

THIS DOCUMENT IS PROVIDED “AS IS”. ARM PROVIDES NO REPRESENTATIONS AND NO WARRANTIES, EXPRESS, IMPLIED OR STATUTORY, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY, SATISFACTORY QUALITY, NON-INFRINGEMENT OR FITNESS FOR A PARTICULAR PURPOSE WITH RESPECT TO THE DOCUMENT. For the avoidance of doubt, Arm makes no representation with respect to, and has undertaken no analysis to identify or understand the scope and content of, patents, copyrights, trade secrets, or other rights.

This document may include technical inaccuracies or typographical errors.

TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL ARM BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF ARM HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

This document consists solely of commercial items. You shall be responsible for ensuring that any use, duplication or disclosure of this document complies fully with any relevant export laws and regulations to assure that this document or any portion thereof is not exported, directly or indirectly, in violation of such export laws. Use of the word “partner” in reference to Arm’s customers is not intended to create or refer to any partnership relationship with any other company. Arm may make changes to this document at any time and without notice.

This document may be translated into other languages for convenience, and you agree that if there is any conflict between the English version of this document and any translation, the terms of the English version of the Agreement shall prevail.

The Arm corporate logo and words marked with ® or ™ are registered trademarks or trademarks of Arm Limited (or its affiliates) in the US and/or elsewhere. All rights reserved. Other brands and names mentioned in this document may be the trademarks of their respective owners. Please follow Arm’s trademark usage guidelines at <http://www.arm.com/company/policies/trademarks>

Copyright © 2021 Arm Limited (or its affiliates). All rights reserved.

Arm Limited. Company 02557590 registered in England.

110 Fulbourn Road, Cambridge, England CB1 9NJ.

LES-PRE-20349 version 21.0

Contents

Preview of Rules-based Scalable Vector Extension (SVE) Supplement

	Release information	ii
	Non-Confidential Proprietary Notice	iii
Preface		
	About this book	viii
	Conventions	ix
	Typographical conventions	ix
	Numbers	ix
	Pseudocode descriptions	ix
	Assembler syntax descriptions	ix
	Rules-based writing	x
	Additional reading	xi
	Arm publications	xi
	Feedback	xii
	Feedback on this book	xii
	Progressive Terminology Commitment	xii
Chapter 1	Introduction	
	1.1 About the SVE supplement	13
	1.2 About the Scalable Vector Extension	14
	1.2.1 Features within SVE	14
	1.3 Half-precision floating-point support	15
	1.4 Register disambiguation	16
Chapter 2	SVE Application level programmers' model	
	2.1 SVE-specific registers	17
	2.1.1 SVE Vector registers	17
	2.1.2 SVE predicate registers	18
	2.1.3 First Fault Register, FFR	19
	2.1.4 SVE writes to scalar registers	20
	2.2 Process state, PSTATE N, Z, C and V Condition flags	21
Chapter 3	Predication	
	3.1 Governing predicate	23
Chapter 4	SVE System level programmers' model	
	4.1 Exception model	24
	4.2 Synchronous memory faults	26
	4.2.1 SVE modifications to precise exceptions and data aborts	26
	4.2.2 SVE asynchronous exception behavior	28
	4.3 Configurable vector length	29
Chapter 5	SVE Memory Model	
	5.1 Atomicity	31
	5.2 Alignment support	32
	5.3 Endian support	33
	5.4 Memory ordering	34
	5.5 Device memory	35

5.6	CONSTRAINED UNPREDICTABLE memory accesses	36
Chapter 6	SVE instruction set	
6.1	SVE assembler language	37
6.2	SVE ISA functional groups	39
6.2.1	Load, store, and prefetch instructions	39
6.2.2	Vector move operations	44
6.2.3	Integer operations	44
6.2.4	Bitwise operations	49
6.2.5	Floating-point operations	50
6.2.6	Predicate operations	57
6.2.7	Move operations	64
6.2.8	Reduction operations	67
Chapter 7	SVE System and control registers	
7.1	SVE modifications to the AArch64 System registers	69
7.2	SVE-specific AArch64 System registers or fields	71
Chapter 8	SVE Debug	
8.1	Self-hosted debug	72
8.1.1	SVE Watchpoint exceptions	72
8.1.2	MOVPRFX instruction behavior in self-hosted debug	73
8.2	External debug	74
8.2.1	Instructions in Debug state	74
Chapter 9	SVE Performance Monitors Extension	
9.1	SVE and PMU interaction	75
9.2	New performance monitor events	77
9.2.1	SVE-specific PMU event descriptions	77
9.3	PMU events	78
Chapter 10	Recommended SVE-specific PMU events	
10.1	Interesting combinations of SVE events	80
10.1.1	Scalar-equivalent operations	80
10.1.2	Bytes loaded and stored	80
10.1.3	Overall vector utilization	80
10.1.4	Vector loop efficiency	81
Chapter 11	Instruction categories	
11.1	Data movement instructions	82
11.1.1	Data movement (scalar)	82
11.1.2	Data movement (Advanced SIMD)	82
11.1.3	Data movement (SVE)	83
11.2	Integer instructions	84
11.2.1	Integer (scalar)	84
11.2.2	Integer (Advanced SIMD)	85
11.2.3	Integer (SVE)	89
11.3	Floating-point instructions	92
11.3.1	Floating-point (scalar)	92
11.3.2	Floating-point (Advanced SIMD)	92
11.3.3	Floating-point (SVE)	94
11.4	Floating-point conversions	96
11.4.1	Float↔Float convert (scalar)	96
11.4.2	Float↔Float convert (Advanced SIMD)	96
11.4.3	Float↔Float convert (SVE)	96
11.4.4	Float↔Int convert (scalar)	96

11.4.5	Float↔Int convert (Advanced SIMD)	96
11.4.6	Float↔Int convert (SVE)	97
11.5	Floating-point or integer instructions	98
11.5.1	Floating-point or integer arithmetic (scalar)	98
11.5.2	Floating-point or integer arithmetic (Advanced SIMD)	98
11.5.3	Floating-point or integer arithmetic (SVE)	98
11.6	Non-SIMD SVE instructions	99
11.6.1	Element count and increment scalar (SVE)	99
11.6.2	Compare and terminate (SVE)	99
11.7	Predicate instructions	100
11.7.1	Predicate move (SVE)	100
11.7.2	Predicate counted loop (SVE)	100
11.7.3	Predicate bitwise logical operations (SVE)	100
11.7.4	Predicate scan (SVE)	100
11.7.5	Predicate count and increment scalar (SVE)	101
11.7.6	Predicate count and increment vector (SVE)	101
11.8	Cryptographic instructions	102
11.8.1	Cryptographic (Advanced SIMD)	102
11.9	Load/store/prefetch instructions	103
11.9.1	Load/store (Advanced SIMD and floating-point scalar)	103
11.9.2	Load/store/prefetch (SVE)	104

Chapter 12

Glossary

Draft

Preface

Draft

About this book

This book is the Arm[®] Architecture Reference Manual Supplement, The Scalable Vector Extension (SVE). This book describes the changes and additions to the ARMv8-A AArch64 architecture that are introduced by SVE, and therefore must be read in conjunction with the ARM[®] Architecture Reference Manual, ARMv8-A, for ARMv8-A architecture profile.

- SVE Application level programmers' model
Describes how the PE at an application level is altered by the implementation of SVE.
- Predication
Describes the relationship between predication, instructions, registers and SVE.
- SVE System level programmer's model
Describes how the PE at a system level is altered by the implementation of SVE.
- SVE Memory Model
Describes the extensions made for SVE to the Arm memory model.
- SVE instruction set
Describes the extensions made for SVE to the Arm instruction set.
- SVE System and control registers
Describes the extensions made for SVE to the Arm System and control registers.
- SVE Debug
Describes how the Arm v8Debug exception model has been extended for SVE.
- SVE Performance Monitors Extension
Describes the interaction between SVE and PMU.
- Recommended SVE-specific PMU events.
Lists the recommended PMU events for SVE implementations
- Instruction categories
Lists SVE instructions grouped by the type of instruction.

Conventions

Typographical conventions

The typographical conventions are:

italic

Introduces special terminology, and denotes citations.

bold

Denotes signal names, and is used for terms in descriptive lists, where appropriate.

`monospace`

Used for assembler syntax descriptions, pseudocode, and source code examples.

Also used in the main text for instruction mnemonics and for references to other items appearing in assembler syntax descriptions, pseudocode, and source code examples.

SMALL CAPITALS

Used for some common terms such as IMPLEMENTATION DEFINED.

Used for a few terms that have specific technical meanings, and are included in the Glossary.

Red text

Indicates an open issue.

Blue text

Indicates a link. This can be

- A cross-reference to another location within the document
- A URL, for example <http://developer.arm.com>

Numbers

Numbers are normally written in decimal. Binary numbers are preceded by `0b`, and hexadecimal numbers by `0x`. In both cases, the prefix and the associated value are written in a monospace font, for example `0xFFFF0000`. To improve readability, long numbers can be written with an underscore separator between every four characters, for example `0xFFFF_0000_0000_0000`. Ignore any underscores when interpreting the value of a number.

Pseudocode descriptions

This book uses a form of pseudocode to provide precise descriptions of the specified functionality. This pseudocode is written in a monospace font. The pseudocode language is described in the Arm Architecture Reference Manual.

Assembler syntax descriptions

This book contains numerous syntax descriptions for assembler instructions and for components of assembler instructions. These are shown in a `monospace` font.

Rules-based writing

Rules-based writing differs significantly from the traditional style used in Arm technical documents. Rules-based writing has short statements with unique identifiers.

Requirements are referred to as rules and other information is referred to as information statements. Rules-based architectural documents describe requirements of the architecture and information statements provide additional information. Structured rules also follow a specific structure and specify some keyword terms. The following is an example of a rule, followed by an information statement.

R_{WPBT}

If a document has structured rules, by default all rules statements have a specific structure.

I_{NBPP}

For architectural specifications the writing style for outcomes is deliberately different from the usual Arm style.

All rules and information statements have unique IDs. IDs start with a designator, followed by a unique string of 4 or 5 SMALL CAPITAL consonants. If the designator is an 'R' it is a rule. If the designator is an 'I', it is an information statement.

All rules have a specific structure. Information statements may take any structure.

Rules generally start with any conditions that make the rule applicable. These conditions have a limited set of introductory phrases:

- Conditions that begin with *if* are used to make rules conditional on a state and tend to last for a while.
- Conditions that begin with *when* are used to make rules conditional on an event happening.
- Conditions that begin with *for* are used to make a rule apply to a part of the system.

There are three forms of the *if*-statements:

- *If* indicates the condition is sufficient to cause the action but might not be necessary. "If X then Y" means the same as "If X then Y happens, but if not X then Y might still happen".
- *Only if* indicates the condition is necessary but might not be sufficient. "Only if X then Y" means the same as "If X then Y might happen, if not X then Y cannot happen".
- *If and only if* indicates the condition is both necessary and sufficient. "If and only if X then Y" means the same as "If X then Y happens, if not X then Y cannot happen".

There are also three forms of the *when*-statements:

- *When* indicates the condition is sufficient to cause the action but might not be necessary.
- *Only when* indicates the condition is necessary but might not be sufficient.
- *When and only when* indicates the condition is both necessary and sufficient.

There may be many preconditions in a rule.

The next part of a rule is either an actor or a subject. When a specific action by a specific entity is defined, the rule will be written in active voice and will have an actor. If the action is performed by an IMPLEMENTATION DEFINED entity, then the rule will be written in passive voice and the rule will have a subject, which is something that is acted on by the action in the rule statement. The action to be carried out follows the actor or subject and is required. The object(s) of the action, followed by the outcome of the rule, are both optional and may be present after the action.

Additional reading

This section lists publications by Arm and by third parties.

See Arm Developer (<http://developer.arm.com>) for access to Arm documentation.

Arm publications

- *Arm[®] Architecture Reference Manual, Armv8-A, for Armv8-A architecture profile (ARM DDI 0487).*
- *System Register XML for Armv8.7.*
- *A64 ISA XML for Armv8.7.*

Draft

Feedback

Arm welcomes feedback on its documentation.

Feedback on this book

If you have comments on the content of this book, send an e-mail to errata@arm.com. Give:

- The title (Preview of Rules-based Scalable Vector Extension (SVE) Supplement).
- The number (AES0047 Aa).
- The page numbers to which your comments apply.
- The rule identifiers to which your comments apply, if applicable.
- A concise explanation of your comments.

Arm also welcomes general suggestions for additions and improvements.

Note

Arm tests PDFs only in Adobe Acrobat and Acrobat Reader, and cannot guarantee the appearance or behavior of any document when viewed with any other PDF reader.

Progressive Terminology Commitment

Arm values inclusive communities. Arm recognizes that we and our industry have used terms that can be offensive. Arm strives to lead the industry and create change.

We believe that this document contains no offensive terms. If you find offensive terms in this document, please contact terms@arm.com

Chapter 1

Introduction

Draft

1.1 About the SVE supplement

- \perp_{BFDMT} This supplement describes the architectural changes that are introduced by the *Scalable Vector Extension (SVE)*
- \perp_{XCFJS} This supplement is to be read in conjunction with the Arm® Architecture Reference Manual and XML for the Armv8-A architecture profile. This provides a comprehensive description of the Armv8-A architecture, including *SVE*.
- \perp_{MXZMC} This supplement does not contain any detailed instruction descriptions, pseudocode, XML, or System register descriptions. This information is provided in a separate format. Links to this information are included throughout the supplement.

1.2 About the Scalable Vector Extension

I_{TKDFM}

The *Scalable Vector Extension* (SVE) includes the following functionality:

- Configurable vector length, from 128 bits up to 2048 bits.
- Predication and the required predicate registers.
- Instructions that operate on variable size vectors and predicates.
- Gather-load and scatter-store.
- Support for software-managed speculative vectorization.
- Some minor additions to the configuration and identification registers.

I_{VKNHX}

SVE complements the AArch64 *Advanced SIMD and floating-point* functionality. *SVE* does not replace the AArch64 *Advanced SIMD and floating-point* functionality.

R_{XZBNC}

If *SVE* is implemented, then FEAT-FP16 and FEAT_FCMA-CompNum are implemented.

R_{JLGXX}

SVE is supported in AArch64 state only.

1.2.1 Features within SVE

R_{YKCSW}

The following list summarizes when specific features within *SVE* are OPTIONAL or mandatory:

- FEAT_BF16 BFloat16 instructions are OPTIONAL in Armv8.2 implementations and mandatory in Armv8.6 implementations. If *SVE* and AArch64 Advanced SIMD are both implemented, they must agree on the presence of FEAT_BF16.
- FEAT_I8MM matrix multiplication instructions are OPTIONAL in Armv8.2 implementations and mandatory in Armv8.6 implementations. If *SVE* and AArch64 Advanced SIMD are both implemented, they must agree on the presence of FEAT_I8MM.
- FEAT_F32MM and FEAT_F64MM matrix multiplication instructions are OPTIONAL for *SVE* in Armv8.2.

1.3 Half-precision floating-point support

R _{MDHHJ}	<p>If the <i>Effective value</i> of FPCR.FZ16 is 0, the <i>SVE</i> half-precision instructions use the same floating-point exception traps and floating-point exception enables as the equivalent <i>SVE</i> single-precision or <i>SVE</i> double-precision instructions.</p> <p>If FPCR.FZ16 is 1, all of the following are true:</p> <ul style="list-style-type: none">• For all of the <i>SVE</i> half-precision instructions, flushing to zero of denormalized inputs is enabled.• For <i>SVE</i> conversion instructions between half-precision and single or double-precision, flushing to zero of denormalized inputs is not enabled.
R _{FCPTB}	<p>If FPCR.AH is 1, <i>SVE</i> instructions flush denormalized numbers to zero following the same behavior as the equivalent scalar floating-point instructions.</p>
R _{DHDYN}	<p>The behavior of <i>SVE</i> half-precision instructions is not affected by the value of the FPCR.FZ bit.</p>
R _{HYPXG}	<p>When a half-precision value is flushed to zero because FPCR.FZ16 is 1, an input denormal exception that sets FPSR.IDC to 1 does not occur.</p>
R _{DLWTQ}	<p>The <i>SVE</i> half-precision floating-point instructions support only the IEEE 754-2008 half-precision format.</p>
R _{RWYCB}	<p><i>SVE</i> half-precision floating-point instructions ignore the value of the FPCR.AHP bit and behave as if the bit has an <i>Effective value</i> of 0.</p>

Draft

1.4 Register disambiguation

I_{LJ}SKW

In some sections of this manual, registers are referred to by a generic name, where the description applies to more than one context. This is because the description applies to multiple *Exception levels*, and therefore at a particular *Exception level* the register names need to take the appropriate *Exception level* suffix, *_EL0*, *_EL1*, *_EL2*, or *_EL3*. The following table disambiguates the generic names of some *System registers* by *Exception level*:

Generic form	EL0	EL1	EL2	EL3
ELR_ELx	-	ELR_EL1	ELR_EL2	ELR_EL3
ESR_ELx	-	ESR_EL1	ESR_EL2	ESR_EL3
FAR_ELx	-	FAR_EL1	FAR_EL2	FAR_EL3
SCTLR_ELx	-	SCTLR_EL1	SCTLR_EL2	SCTLR_EL3
ZCR_ELx	-	ZCR_EL1	ZCR_EL2	ZCR_EL3

Draft


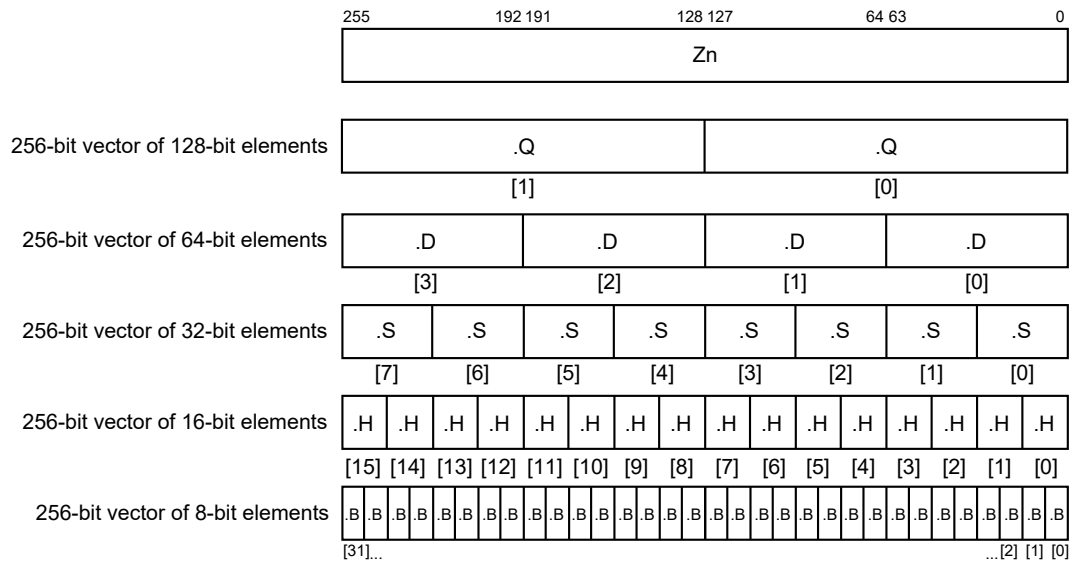
Chapter 2

SVE Application level programmers' model

2.1 SVE-specific registers

2.1.1 SVE Vector registers

R _{FBSJ}	SVE has 32 scalable vector registers named Z0-Z31.
R _{WNJD}	All SVE scalable vector registers are the same size.
R _{KCQB}	The size of an SVE scalable vector register is an IMPLEMENTATION DEFINED multiple of 128 bits.
R _{KSDQ}	The maximum size of an SVE scalable vector register is 2048 bits.
R _{RPHX}	The minimum size of an SVE scalable vector register is 128 bits.
I _{GKJYJ}	Unless stated otherwise in an instruction description, SVE instructions treat an SVE scalable vector register as containing one or more vector elements that are equal in size.
I _{CDKJQ}	Unless stated otherwise in an instruction description, vector elements can be processed in parallel by SVE instructions.
R _{KHBN}	When an SVE scalable vector register is divided into vector elements by an instruction, the size of the vector elements is encoded in the opcode of the instruction. The size of the vector elements is 8, 16, 32, 64, or 128 bits.
R _{CJZLM}	When the order of operations performed by an SVE instruction on vector or predicate elements has observable significance, elements are processed in increasing element number order.
I _{DBZRX}	The layouts of an SVE 256-bit vector register and a SIMD&FP vector in AArch64 state are:



R_{YDXCP} Bits[127:0] of each of the SVE scalable vector registers, Z0-Z31, hold the correspondingly numbered AArch64 SIMD&FP register, V0-V31.

R_{WKYLB} When the accessible SVE vector length at the current Exception level is greater than 128 bits, any AArch64 instruction that writes to V0-V31 sets all the accessible bits above bit [127] of the corresponding SVE scalable vector register to zero.

See also:

- [2.1.2 SVE predicate registers.](#)
- [4.3 Configurable vector length.](#)
- [5.1 Atomicity.](#)

2.1.2 SVE predicate registers

R_{DCWFB} SVE has 16 scalable predicate registers named P0-P15.

R_{NLGZS} Each SVE predicate register holds one bit for each byte of a vector register.

R_{NKRJV} The size of an SVE predicate register is an IMPLEMENTATION DEFINED multiple of 16 bits.

Chapter 2. SVE Application level programmers' model

2.1. SVE-specific registers

R _{MFPXG}	The maximum size of an SVE predicate register is 256 bits.
R _{BBTXX}	The minimum size of an SVE predicate register is 16 bits.
R _{XVRRX}	Unless stated otherwise in the instruction description, SVE instructions treat an SVE predicate register as containing one or more predicate elements of equal size.
R _{XMPLM}	Each predicate register can be subdivided into a number of 1, 2, 4, or 8-bit elements.
R _{NSXCV}	Each predicate element in a predicate register corresponds to a vector element.
R _{XCZQR}	When a predicate register is divided into predicate elements by an instruction, the size of the predicate elements is encoded in the opcode of the instruction.
R _{DNMFH}	If the lowest-numbered bit of a predicate element is 1, the value of the predicate element is TRUE.
R _{HRPMD}	If the lowest-numbered bit of a predicate element is 0, the value of the predicate element is FALSE.
R _{HBMLS}	For all SVE instructions, if all of the following are true, all bits except the lowest-numbered bit of each predicate element are ignored on reads: <ul style="list-style-type: none">• The instructions are not used to move and permute predicate elements.• The instructions are not predicate logical operations.
R _{LTGQC}	For all SVE instructions, if all of the following are true, all bits except the lowest-numbered bit of each predicate element are set to zero on writes: <ul style="list-style-type: none">• The instructions are not used to move and permute predicate elements.• The instructions are not predicate logical operations.

See also:

- [Chapter 3 Predication](#).
- [6.2.6.3 Predicate logical operations](#) This section contains a list of instructions used to perform bitwise logical operations on predicate registers that operate on all bits of the register.
- [6.2.7.3 Predicate permute](#) This section contains a list of instructions used to used to move and permute predicate elements.
- [11.7 Predicate instructions](#).

2.1.3 First Fault Register, FFR

R _{TRLWH}	SVE has a dedicated First Fault Register named FFR.
I _{XPLQW}	The FFR captures the cumulative fault status of a sequence of SVE First-fault and Non-fault vector load instructions.
R _{CPQQN}	The FFR and the predicate registers have the same size and format.
I _{PBWPB}	The FFR is a Special-purpose register.
R _{CGHCK}	All accessible bits in the FFR are initialized to 1 by using the SETFFR instruction.
R _{WZJVT}	Bits in the FFR are indirectly set to 0 as a result of a suppressed access or fault generated in response to an <i>Active element</i> of an SVE First-fault or Non-fault vector load.
R _{BZLJG}	Bits in the FFR are never set to 1 as a result of a vector load instruction.
I _{XLZQY}	After a sequence of one or more SVE First-fault or Non-fault loads that follow a SETFFR instruction, the FFR contains a sequence of zero or more TRUE elements, followed by zero or more FALSE elements.
I _{TQMTV}	The TRUE elements in the FFR indicate the shortest sequence of consecutive elements that could contain valid data loaded from memory.
R _{GHRQ}	The only instructions that directly read the FFR are: <ul style="list-style-type: none">• R_{DFFR}.

- R_DFFRS.

R_{LHBRN} The only instructions that directly write the FFR are:

- WRFFR.
- SETFFR.

R_{XXMMP} All direct and indirect reads and writes to the FFR occur in program order relative to other instructions, without explicit synchronization.

See also:

- [2.1.2 SVE predicate registers.](#)
- [4.2 Synchronous memory faults.](#)

2.1.4 SVE writes to scalar registers

I_{ZDLGD} Certain SVE instructions generate a scalar result that is written to an AArch64 general-purpose register or to element[0] of a vector register.

R_{HNVTM} When an SVE instruction generates a scalar result of width N bits, the instruction places the result in bits [N-1:0] of the destination register.

R_{QCLSH} When an instruction generates a scalar result of width N bits, and N is less than the maximum destination register width RW, the instruction sets to zero bits [RW-1:N] of the destination register.

See also:

- *Registers in AArch64 Execution state* in the ARM[®] Architecture Reference Manual, ARMv8-A, for ARMv8-A architecture profile

2.2 Process state, PSTATE N, Z, C and V Condition flags

I_{WYXLS} Process state, or *PSTATE*, is an abstraction of process state information. This section describes the *SVE*-specific use of *PSTATE*.

I_{YZYCQ} PSTATE N, Z, C and V Condition flags can be updated by any of the following:

- An *SVE* instruction that generates a predicate result and updates the PSTATE N, Z, C and V Condition flags based on the value of the result.
- An *SVE* instruction that updates the PSTATE N, Z, C and V Condition flags based on the value in its predicate source register or FFR:
 - PTEST
 - RDIFFRS (predicated)
- An *SVE* instruction that updates the PSTATE N, Z, C and V Condition flags based on the values in its general-purpose source registers:
 - CTERM EQ
 - CTERM NE

R_{TPXTF} When setting the PSTATE N, Z, C and V Condition flags for *SVE* predicated flag-setting instructions, the instruction's Governing predicate determines which predicate elements are considered Active.

R_{QJBRW} When setting the PSTATE N, Z, C and V Condition flags for *SVE* unpredicated flag-setting instructions, all predicate elements are considered Active.

R_{ZMRXC} Unless otherwise specified in an instruction description, the *SVE* flag-setting instructions update the PSTATE N, Z, C and V Condition flags as follows:

Flag	SVE Name	SVE interpretation
N	First	Set to 1 if the <i>First active element</i> was TRUE, otherwise cleared to 0.
Z	None	Cleared to 0 if any <i>Active element</i> was TRUE, otherwise set to 1.
C	Not last	Cleared to 0 if the <i>Last active element</i> was TRUE, otherwise set to 1.
V	-	Cleared to 0.

I_{KSXVR} For convenience, the *SVE* assembler syntax defines an alternative set of *SVE* condition code aliases for use with AArch64 conditional instructions, as follows:

Condition test	AArch64 name	SVE alias	SVE interpretation
Z == 1	EQ	NONE	All <i>Active elements</i> were FALSE or there were no <i>Active elements</i>
Z == 0	NE	ANY	An <i>Active element</i> was TRUE.
C == 1	HS/CS	NLAST	The <i>Last active element</i> was FALSE.
C == 0	LO/CC	LAST	The <i>Last active element</i> was TRUE.
N == 1	MI	FIRST	The <i>First active element</i> was TRUE.
N == 0	PL	NFRST	The <i>First active element</i> was FALSE.
C == 1 && Z == 0	HI	PMORE	An <i>Active element</i> was TRUE, but the <i>Last active element</i> was FALSE.

Condition test	AArch64 name	SVE alias	SVE interpretation
$C == 0 \parallel Z == 1$	LS	PLAST	The <i>Last active element</i> was TRUE, or all <i>Active elements</i> were FALSE, or there were no <i>Active elements</i> .
$V == 1$	VS	-	CTERM comparison failed, but end of partition reached.
$V == 0$	VC	-	CTERM comparison succeeded, or end of partition not reached.
$N == V$	GE	TCONT	CTERM termination condition not detected.
$N != V$	LT	TSTOP	CTERM termination condition detected.

See also:

- [3.1 Governing predicate](#)

Draft

Chapter 3

Predication

Draft

3.1 Governing predicate

<code>I_{KNKBN}</code>	Many predicated instructions can only use P0-P7 as the Governing predicate.
<code>I_{HGKJY}</code>	For more information on predicated instructions, see the individual instruction descriptions.
<code>R_{LZVFJ}</code>	When a Governing predicate element is TRUE, the corresponding vector or predicate register elements are Active.
<code>R_{CNFLG}</code>	When a Governing predicate element is FALSE, the corresponding vector or predicate register elements are Inactive.
<code>R_{CBYJH}</code>	Predicated instructions process <i>Active elements</i> .
<code>R_{LDXSF}</code>	Predicated instructions do not process <i>Inactive elements</i> .
<code>R_{WLQBD}</code>	When a predicated instruction writes to a vector destination register or a predicate destination register, one of the following happens: <ul style="list-style-type: none">• The <i>Inactive elements</i> in the destination register are set to zero.• The <i>Inactive elements</i> in the destination register retain their previous value.
<code>I_{QBHRN}</code>	Zeroing predication is performed when the <i>Inactive elements</i> in the destination register are set to zero.
<code>I_{YPYRF}</code>	Merging predication is performed when <i>Inactive elements</i> in the destination register retain their previous value.
<code>R_{DSZRY}</code>	When an unpredicated SVE flag-setting instruction sets flags, all elements in the predicate source and destination registers are treated as Active.

See also:

- [11.7 Predicate instructions](#) Lists the predicated instructions.

Chapter 4

SVE System level programmers' model

Draft

4.1 Exception model

I _{QNHYT}	<p>SVE adds hierarchical trap and enable controls at EL3, EL2, and EL1:</p> <ul style="list-style-type: none">• CPTR_EL3.EZ.• CPTR_EL2.TZ, when HCR_EL2.E2H == 0.• CPTR_EL2.ZEN, when HCR_EL2.E2H == 1.• CPACR_EL1.ZEN.
I _{FYMP}	<p>SVE defines the 0b011001 exception class value in ESR_ELx.EC. The 0b011001 exception class value is for exceptions that are due to attempted execution of SVE instructions and MRS/MSR accesses to the ZCR_ELx registers that are trapped by CPACR_EL1, CPTR_EL2 or CPTR_EL3.</p>
I _{BJRCZ}	<p>The 0b011001 exception class value is for exceptions that are due to attempted execution of SVE instructions and MRS or MSR accesses to the ZCR_ELx registers that are trapped by CPACR_EL1[./SVE_SysReg/xhtml/CPACR_EL1.html], CPTR_EL2 or CPTR_EL3.</p>
R _{VQCKK}	<p>SVE floating-point instructions only generate floating-point exceptions in response to floating-point operations performed on <i>Active elements</i>.</p>
R _{RWVTR}	<p>When a MOVPRFX instruction pairs legally with another instruction and the execution of the pair generates a synchronous exception, the return address that is stored in ELR_ELx is one of the following:</p> <ul style="list-style-type: none">• When the MOVPRFX instruction did not cause a change to the architectural state, the address of the MOVPRFX instruction is stored.• When the MOVPRFX instruction caused a change to the architectural state, the address of the prefixed instruction is stored.

R_{XRWVD}	<p>When a MOVPRFX instruction pairs legally with another instruction and the execution of the pair causes entry to <i>Debug state</i>, the return address that is stored in DLR_EL0 is one of the following:</p> <ul style="list-style-type: none"> • When the MOVPRFX instruction did not cause a change to the architectural state, the address of the MOVPRFX instruction is stored. • When the MOVPRFX instruction caused a change to the architectural state, the address of the prefixed instruction is stored.
R_{TPRKM}	<p>When a MOVPRFX instruction pairs illegally with another instruction and execution of the pair generates a synchronous exception, the return address recorded in ELR_ELx is a CONSTRAINED UNPREDICTABLE choice one of the following:</p> <ul style="list-style-type: none"> • The address of the MOVPRFX instruction. • The address of the prefixed instruction.
R_{JVNGC}	<p>When a MOVPRFX instruction pairs illegally with another instruction and execution of the pair causes entry to <i>Debug state</i>, the return address recorded in DLR_EL0 is a CONSTRAINED UNPREDICTABLE choice one of the following:</p> <ul style="list-style-type: none"> • The address of the MOVPRFX instruction. • The address of the prefixed instruction.
R_{CRRPM}	<p>When a prefixed instruction generates an Instruction Abort due to an MMU fault or synchronous <i>External abort</i> and the MOVPRFX does not generate an Instruction Abort, then the address of the prefixed instruction is recorded in the appropriate FAR_ELx or HPFAR_EL2 register and the address of the MOVPRFX instruction is recorded in the appropriate ELR_ELx register.</p>
R_{ZJYDX}	<p>When a prefixed instruction generates an Instruction Abort due to an MMU fault or synchronous <i>External abort</i> and the MOVPRFX also generates an Instruction Abort, then the address of the MOVPRFX instruction is recorded in the appropriate FAR_ELx or HPFAR_EL2 register and the appropriate ELR_ELx register.</p>

See also:

- ESR_EL1
- ESR_EL2
- ESR_EL3
- ELR_EL1
- ELR_EL2
- ELR_EL3
- FAR_EL1
- FAR_EL2
- FAR_EL3
- MOVPRFX (predicated)
- MOVPRFX (unpredicated)
- ZCR_EL1
- ZCR_EL2
- ZCR_EL3
- [6.2.7.5 Move prefix](#)

4.2 Synchronous memory faults

I _{DGSNC}	<p>SVE load and store instructions can generate a memory access sequence that might not be completed due to an exception occurring during the memory access sequence.</p>
R _{PGRKJ}	<p>When a memory access resulting from an SVE load or store instruction causes a synchronous memory fault, all of the following occur:</p> <ul style="list-style-type: none">• The ESR_ELx.EC field is set to one of the following values:<ul style="list-style-type: none">– If the synchronous memory fault is taken to a higher <i>Exception level</i>, the ESR_ELx.EC field is set to 0b100100.– If the synchronous memory fault is taken to the same <i>Exception level</i>, the ESR_ELx.EC field is set to 0b100101.• The appropriate fault address register is updated with the lowest address that applies to the <i>Active element</i> for which the fault was reported.<ul style="list-style-type: none">– For stage 1 synchronous memory faults, FAR_ELx at the target <i>Exception level</i> is updated with the fault address.– For stage 2 synchronous memory faults, HPFAR_EL2 is updated with the fault address.• A Data Abort exception is then taken.
R _{ZKBRX}	<p>When the execution of an SVE instruction causes multiple faults from different memory addresses, the multiple faults are not prioritized by the architecture.</p>
R _{KWRHM}	<p>When returning from an exception, an SVE load or store instruction that has not been architecturally executed is restarted from the beginning of its processing.</p>
I _{MSVYK}	<p>When an SVE load or store instruction that has not been architecturally executed is restarted due to returning from an exception, one or more of the memory locations might have been accessed multiple times. Therefore, there might be multiple accesses to memory locations that have been changed between the accesses, even if the memory locations are sensitive to the number of accesses.</p>
R _{SKNTR}	<p>When an SVE load or store instruction results in a data access, any of the following are considered faults:</p> <ul style="list-style-type: none">• MMU fault.• Alignment fault, excluding the SP alignment fault.• Synchronous External abort, including synchronous parity error or ECC error.• Watchpoint debug event. <p>See also:</p> <ul style="list-style-type: none">• ESR_EL1• ESR_EL2• ESR_EL3• FAR_EL1• FAR_EL2• FAR_EL3

4.2.1 SVE modifications to precise exceptions and data aborts

I _{BSRZK}	<p>For SVE vector load and store instructions that generate a fault, SVE modifies the ARMv8 A architecture profile's standard behaviors for dealing with precise exceptions and data aborts. The rules in this section define the SVE specific behaviors.</p>
I _{FFMJQ}	<p>In this section, the term fault refers to any item listed in R-SKNTR.</p>
R _{LXNVS}	<p>When the memory access due to the <i>First active element</i> of an SVE First-fault vector load instruction generates a fault, the <i>First active element</i> is handled the same way as any <i>Active element</i> of a predicated SVE vector load instruction.</p>

- R_{BZTBN}** When the memory access due to the *First active element* of an SVE First-fault vector load instruction does not generate a fault, the other elements are handled in the same way as the elements of an SVE Non-fault vector load instruction.
- R_{DWYCY}** When an SVE predicated vector store causes a fault, one or more of the following occurs:
- Memory locations that are associated with *Active elements* that do not generate a fault are set to an UNKNOWN value.
 - Memory locations that are associated with *Inactive elements* are preserved.
 - Memory locations that are associated with *Active elements* that generate a fault are preserved.
- R_{ZXNXT}** When execution of an SVE load instruction causes a fault and the destination is not a vector register that is also used as a base or index register by the instruction, then all elements of the destination register will contain an UNKNOWN value.
- R_{SNJQR}** When execution of an SVE load instruction causes a fault and the destination is a vector register that is also used as a base or index register by the instruction, then all elements of the destination register will retain their original value prior to execution of the load instruction.
- R_{JKGYJ}** When a fault is detected for any of the following, an exception is not taken and the ESR_ELx exception syndrome, FAR_ELx, and HPFAR_EL2 fault address registers are not updated.
- Any *Active element* of an SVE Non-fault vector load.
 - Any *Active element* of an SVE First-fault vector load except for the *First active element*.
- R_{MNKNV}** The PE can suppress the read of any *Active element* of an SVE Non-fault vector load or any *Active element* from a First-fault vector load except for the *First active element*.
- R_{YFTRN}** When any of the following memory accesses generates a fault or is suppressed for any other reason, the FFR predicate elements starting from that element number, up to and including the highest-numbered element, are set to FALSE:
- Any element of a Non-fault vector load.
 - Any element from a First-fault vector load except for the *First active element*.
- I_{BMQVT}** An FFR predicate element is never set to TRUE by an SVE vector load, therefore the fault indications are cumulative.
- R_{NGFTJ}** After an SVE Non-fault vector load is executed, each destination vector element contains one of the values listed in the following table.
- After an SVE First-fault vector load is executed, each destination vector element except for the *First active element* contains one of the values listed in the following table.

Corresponding FFR element	Vector element status	Content of destination vector element
FALSE	Active	Each byte of the element contains an independently CONSTRAINED UNPREDICTABLE choice of one of the following: <ul style="list-style-type: none"> • 0. • The previous value of that byte in the destination vector register. • If and only if all of the following apply, the value read from memory: <ul style="list-style-type: none"> – The memory access for that byte was not an access to any type of Device memory. – The memory access for that byte does not return information that cannot be accessed at the current or a lower level of privilege.

Corresponding FFR element	Vector element status	Content of destination vector element
FALSE	Inactive	A CONSTRAINED UNPREDICTABLE choice of: <ul style="list-style-type: none"> • 0. • The previous value of that vector element.
TRUE	Active	The value read from memory.
TRUE	Inactive	0.

R_{WCHSR} In the previous table, a watchpoint is not a mechanism for preventing access to memory.

See also:

- *Definition of a precise exception in the ARM® Architecture Reference Manual, ARMv8-A, for ARMv8-A architecture profile*
- *Effect of Data_Aborts and Watchpoints in the ARM® Architecture Reference Manual, ARMv8-A, for ARMv8-A architecture profile*
- ESR_EL1
- ESR_EL2
- ESR_EL3
- FAR_EL1
- FAR_EL2
- FAR_EL3

4.2.2 SVE asynchronous exception behavior

R_{TFLTX} Permitting SVE instructions to be interrupted by asynchronous exceptions is IMPLEMENTATION DEFINED.

R_{WFMZK} When returning from an asynchronous exception, an interrupted SVE instruction is restarted and cannot resume at the point the instruction was interrupted.

4.3 Configurable vector length

- I_{NWYBP}** Privileged *Exception levels* can use the `ZCR_ELx.LEN` *System register* fields to constrain the vector length at that *Exception level* and at less privileged *Exception levels*.
- R_{PVRSF}** An implementation allows the vector length to be constrained to any power of two that is less than the maximum implemented vector length.
- R_{RYQYY}** An implementation is permitted to allow the vector length to be constrained to multiples of 128 that are not a power of two. It is IMPLEMENTATION DEFINED which of the permitted multiples of 128 are supported.
- I_{CPZLW}** The following table shows the *SVE* configurable vector lengths:

Maximum	Required	Permitted
128	128	-
256	128, 256	-
384	128, 256	-
512	128, 256	384
640	128, 256, 512	384
768	128, 256, 512	384, 640
896	128, 256, 512	384, 640, 768
1024	128, 256, 512	384, 640, 768, 896
1152	128, 256, 512, 1024	384, 640, 768, 896
1280	128, 256, 512, 1024	384, 640, 768, 896, 1152
1408	128, 256, 512, 1024	384, 640, 768, 896, 1152, 1280
1536	128, 256, 512, 1024	384, 640, 768, 896, 1152, 1280, 1408
1664	128, 256, 512, 1024	384, 640, 768, 896, 1152, 1280, 1408, 1536
1792	128, 256, 512, 1024	384, 640, 768, 896, 1152, 1280, 1408, 1536, 1664
1920	128, 256, 512, 1024	384, 640, 768, 896, 1152, 1280, 1408, 1536, 1664, 1792
2048	128, 256, 512, 1024	384, 640, 768, 896, 1152, 1280, 1408, 1536, 1664, 1792, 1920

- R_{MMCTJ}** When an unsupported vector length is requested in `ZCR_ELx`, the implementation is required to select the largest supported vector length that is less than the requested vector length. This does not alter the value of `ZCR_ELx.LEN`.
- R_{PXZTM}** If executing at an *Exception level* that is constrained to use a vector length that is less than the maximum implemented vector length, the bits beyond the constrained length of the vector registers, predicate registers, or FFR are inaccessible.
- R_{NLYDK}** When taking an exception from an *Exception level* that is more constrained to a target *Exception level* that is less constrained, the previously inaccessible bits that become accessible have one of the following:
- A value of 0.
 - The value that they had before executing at the more constrained vector size.
- The choice between these options is IMPLEMENTATION DEFINED and can vary dynamically.
- R_{TQVGX}** When the size of the maximum vector length is increased by writing a larger value to `ZCR_ELx.LEN`, the previously inaccessible bits that become accessible have one of the following:

- A value of 0.
- The value that they had before executing at the more constrained vector size.

The choice between these options is IMPLEMENTATION DEFINED and can vary dynamically.

R_{DMBPN}

If both floating-point and SVE instructions are disabled, trapped, or not available at all *Exception levels* below the target *Exception level*, for the current *Security state*, the accessible SVE register state at the target *Exception level* is preserved.

R_{KXKNK}

If SVE instructions are disabled or trapped at EL_x, or not available because that *Exception level* is in AArch32 state, then for all purposes other than a direct read, the ZCR_EL_x.LEN field has an Effective value of 0, which implies an SVE vector length of 128 bits.

See also:

- ZCR_EL1
- ZCR_EL2
- ZCR_EL3

Draft

Chapter 5

SVE Memory Model

Draft

5.1 Atomicity

- I_{HGFJZ}** Atomicity rules for *SIMD* load and store instructions apply to *SVE* load and store instructions. Additional rules apply to atomicity in scalable vector processing.
- I_{CZFSY}** *SVE* predicated memory operations have a vector element size and a memory element size. The vector element size specifies the data that is read from and written to the vector. The memory element size specifies the amount of data that is read from and written to the memory.
- I_{TJQJF}** The vector element size and the memory element size do not need to have the same value.
- I_{LGGHH}** For each memory element, there is an associated element address.
- R_{NNHQB}** *SVE* predicated load and store instructions are performed as a sequence of memory element accesses.
- R_{HLHZV}** When an *SVE* predicated load or store instruction uses an element address that is aligned to the specified memory element size, the related element memory access is performed as a single-copy atomic access.
- R_{JJFRG}** *SVE* unpredicated load and store instructions are performed as a sequence of byte accesses.
- R_{SPYMV}** *SVE* unpredicated load and store instructions do not guarantee that any access larger than a byte will be performed as a single-copy atomic access.

See also:

- Atomicity in the *ARM[®] Architecture Reference Manual, ARMv8-A, for ARMv8-A architecture profile*

5.2 Alignment support

I _{PQQGP}	Alignment rules on <i>SIMD</i> load and store instructions apply to <i>SVE</i> load and store instructions. Additional rules apply to alignment in scalable vector processing.
R _{ZTJZY}	For predicated <i>SVE</i> vector element and structure load or store instructions, alignment checks are based on the memory element size, not on the vector element size.
R _{XFVKZ}	For predicated <i>SVE</i> vector element and structure load or store instructions, <i>Inactive elements</i> cannot cause an <i>Alignment fault</i> .
R _{SCXJP}	For unpredicated <i>SVE</i> vector register load or store instructions, the base address is checked for 16-byte alignment.
R _{CWKLF}	For unpredicated <i>SVE</i> predicate register load or store instructions, the base address is checked for 2-byte alignment.
R _{DLVDL}	If <i>SP</i> alignment checking is enabled in <i>SCTLR_ELx</i> at the current <i>Exception level</i> and an <i>SVE</i> predicated load or store instruction with any <i>Active elements</i> uses the current <i>SP</i> as the base address, then the <i>SP</i> register is checked for 16-byte alignment.
R _{FNCJX}	If <i>SP</i> alignment checking is enabled in <i>SCTLR_ELx</i> at the current <i>Exception level</i> and an <i>SVE</i> predicated load or store instruction with no <i>Active elements</i> uses the current <i>SP</i> as the base address, then it is <i>CONSTRAINED UNPREDICTABLE</i> whether the <i>SP</i> register is checked for 16-byte alignment.

See also:

- *Alignment support* in the *ARM® Architecture Reference Manual, ARMv8-A, for ARMv8-A architecture profile*
- *Memory types and attributes* in the *ARM® Architecture Reference Manual, ARMv8-A, for ARMv8-A architecture profile*
- *SCTLR_EL1*
- *SCTLR_EL2*
- *SCTLR_EL3*
- *SP alignment checking* in the *ARM® Architecture Reference Manual, ARMv8-A, for ARMv8-A architecture profile*
- [11.9 Load/store/prefetch instructions](#)

5.3 Endian support

I _{VDGQK}	Rules on byte and element order of SIMD load and store instructions apply to <i>SVE</i> load and store instructions. Additional rules apply to endianness in scalable vector processing.
R _{CNKCL}	For predicated <i>SVE</i> vector element and structure load and store instructions, an endianness conversion is performed using the size of the memory element. The size of the vector element is not used in endianness conversion.
R _{QHXL}	For unpredicated <i>SVE</i> vector register load and store instructions, the vector byte elements are transferred in increasing element number order without any endianness conversion.
R _{RWLXY}	For unpredicated <i>SVE</i> predicate register load and store instructions, each 8 bits from the predicate are transferred as a byte in increasing element number order without any endianness conversion.
R _{YGSBQ}	When an <i>SVE</i> load instruction is executed, any endianness conversion occurs before any sign-extension or zero-extension into a vector element.
R _{KYRQW}	When an <i>SVE</i> store instruction is executed, any endianness conversion occurs after any truncation from the vector element to the memory element access size.

See also:

- *Data endianness* in the *ARM[®] Architecture Reference Manual, ARMv8-A, for ARMv8-A architecture profile*

Draft

5.4 Memory ordering

<code>I_{CTNGV}</code>	The Armv8 memory model is relaxed for reads and writes generated by <i>SVE</i> load and store instructions.
<code>R_{QLJPC}</code>	When two reads generated by <i>SVE</i> vector load instructions have an address dependency, the dependency does not contribute to the dependency-ordered-before relation.
<code>R_{YMBMZ}</code>	When a pair of reads access the same location, and at least one of the reads is generated by an <i>SVE</i> load instruction, for a given observer, the pair of reads is not required to satisfy the internal visibility requirement.
<code>R_{CJHWV}</code>	When a single <i>SVE</i> vector store instruction generates multiple writes to the same location, the instruction ensures that these writes appear in the coherence order for that location, in order of increasing vector element number. No other ordering restrictions apply to memory effects generated by the same <i>SVE</i> store instruction.
<code>R_{VMDYZ}</code>	If an address dependency exists between two memory reads and an <i>SVE</i> non-temporal vector load instruction generated the second read, then in the absence of any other barrier mechanism to achieve order, the memory accesses can be observed in any order by the other observers within the shareability domain of the memory addresses being accessed.
<code>I_{CCWGN}</code>	For all <i>SVE</i> instructions that load or store one or more vector registers or predicate registers and that generate more than one single-copy atomic access for that load or store, there is no requirement for the memory system beyond the <i>PE</i> to be able to identify the single-copy atomic memory element access sizes.
<code>R_{LVXTJ}</code>	For accesses to different locations of Normal or Device memory, the order in which elements and registers are accessed is not architecturally defined. For multiple writes to the same Normal or Device memory location, the order in which elements and registers are accessed is architecturally defined as increasing from lower to higher vector element numbers.

See also:

- *Definition of the Armv8 memory model, Arm® Architecture Reference Manual*
- *Memory types and attributes, Arm® Architecture Reference Manual*

5.5 Device memory

R _{XLLJZ}	All rules applying to <i>Device memory</i> accesses by <i>Advanced SIMD</i> and floating-point load and store instructions apply to <i>Device memory</i> access by <i>SVE</i> load and store instructions.
I _{YHWJT}	Additional rules apply to <i>Device memory</i> access by <i>SVE</i> load and store instructions.
R _{NYMWH}	When an <i>SVE</i> vector prefetch instruction is executed, any resulting memory read is guaranteed not to access <i>Device memory</i> .
R _{SBHLD}	When an <i>SVE</i> Non-fault vector load is executed or for any element from a First-fault load except the <i>First active element</i> , the resulting memory read will not access <i>Device memory</i> .
R _{TMVNR}	When an <i>SVE</i> Non-fault vector load instruction is executed, an attempt by any <i>Active element</i> to access <i>Device memory</i> is suppressed and reported in the FFR.
R _{SFBKQ}	When an <i>SVE</i> First-fault vector load instruction is executed, any resulting memory read can access <i>Device memory</i> only for the instruction's <i>First active element</i> .
R _{BHNQN}	When an <i>SVE</i> First-fault vector load instruction is executed, an attempt by any <i>Active element</i> other than the <i>First active element</i> to access <i>Device memory</i> is suppressed and is reported in the FFR.
R _{VHCGD}	Hardware speculation of data accesses performed to a <i>Device memory</i> location is not permitted unless stated otherwise.
R _{QHHRG}	For reads, including hardware speculation, that are performed by an <i>SVE</i> unpredicated load instruction, all of the following are true: <ul style="list-style-type: none">• For any 64-byte window aligned to 64 bytes, if at least 1 byte that is explicitly accessed by the instruction, any byte in the window can be accessed by the instruction.• All bytes accessed by the instruction will be in a 64-byte window aligned to 64 bytes containing at least 1 byte that is explicitly accessed by the instruction.
R _{SCGGF}	For reads, including hardware speculation, that are performed by an <i>SVE</i> predicated load instruction that is not a non-temporal load, all of the following are true: <ul style="list-style-type: none">• For any 64-byte window aligned to 64 bytes containing at least 1 byte that is explicitly accessed by an <i>Active element</i> of the instruction, any byte in the window can be accessed by the instruction.• All bytes accessed by the instruction will be in a any 64-byte window aligned to 64 bytes that contains at least 1 byte that is explicitly accessed by an <i>Active element</i> of the instruction.
R _{JWYBD}	For reads, including hardware speculation, that are performed by an <i>SVE</i> predicated non-temporal load instruction from memory locations with the Gathering attributes, all of the following are true: <ul style="list-style-type: none">• For any 128-byte window aligned to 128 bytes containing at least 1 byte that is explicitly accessed by an <i>Active element</i> of the instruction, any byte in the window can be accessed by the instruction.• All bytes accessed by the instruction are in a 128-byte window aligned to 128 bytes that contains at least 1 byte that is explicitly accessed by an <i>Active element</i> of the instruction.
R _{QBLMZ}	Any access to <i>Device memory</i> performed by an <i>SVE</i> load or store instruction is relaxed such that it might behave as if: <ul style="list-style-type: none">• The Gathering attribute is set, regardless of the configured value of the nG attribute.• The Reordering attribute is set, regardless of the configured value of the nR attribute.• The Early Acknowledgement attribute is set, regardless of the configured value of the nE attribute.

Whether or not attributes are classified as mismatched is determined strictly by the memory attributes derived from the page-table entry.

See also:

- [2.1.3 First Fault Register, FFR](#)
- [4.2 Synchronous memory faults](#)

5.6 CONSTRAINED UNPREDICTABLE memory accesses

I _{XZBLN}	CONSTRAINED UNPREDICTABLE behaviors that are associated with memory accesses due to loads and stores also apply to <i>SVE</i> vector load and store instructions.
I _{XFHKS}	The CONSTRAINED UNPREDICTABLE behaviors referred to in this section are defined in the <i>Crossing a page boundary with different memory types or Shareability attributes</i> and <i>Crossing a peripheral boundary with a Device access</i> sections of the the <i>Arm® Architecture Reference Manual, Armv8-A</i> . architecture profile.
R _{SPHTR}	When an <i>SVE</i> unpredicated contiguous load or store instructions accesses an address range that crosses a boundary between memory types, <i>Shareability</i> attributes, or peripherals, the instruction has CONSTRAINED UNPREDICTABLE behaviors associated with the cross boundary memory access.
R _{WPBDN}	When an <i>SVE</i> predicated contiguous load or store instruction has accesses that are associated with <i>Active elements</i> on both sides of a boundary between different memory types, <i>Shareability</i> attributes, or peripherals, the instruction has CONSTRAINED UNPREDICTABLE behaviors associated with the cross boundary memory access.
R _{TGZYT}	When an <i>SVE</i> predicated non-contiguous load or store instruction has a memory access that is associated with an <i>Active element</i> that crosses a boundary between memory types, <i>Shareability</i> attributes, or peripherals, the instruction has CONSTRAINED UNPREDICTABLE behaviors associated with the cross boundary memory access.
R _{FQLTC}	Memory addresses that are associated with <i>Inactive elements</i> cannot trigger CONSTRAINED UNPREDICTABLE behaviors.
R _{FDHZH}	For <i>SVE</i> vector loads and stores that trigger a CONSTRAINED UNPREDICTABLE behavior that then generates an alignment fault, the fault is handled the same as any other synchronous memory fault caused by an <i>SVE</i> load or store instruction.

See also:

- *Crossing a page boundary with different memory types*, *Arm® Architecture Reference Manual, Armv8-A*
- *Crossing a peripheral boundary with a Device access*, *Arm® Architecture Reference Manual, Armv8-A*
- [4.2 Synchronous memory faults](#)

Chapter 6

SVE instruction set

Draft

6.1 SVE assembler language

The *SVE* assembler language extends the A64 assembler language.

I_{RFMVH}

The additions are:

- *SVE* adds vector register names *Z0-Z31* and predicate register names *P0-P15*.
- The number of elements in a vector or predicate register is not specified as part of a vector register shape qualifier. For example, *Z1.B* is used rather than *V1.16B*.
- An element size qualifier is not required for the Governing predicate (*Pg*) except where the element size cannot be inferred from the source and destination element sizes. However, when a predicate element size qualifier is provided it is accepted by an assembler and checked for consistency with the other operands.
- When appropriate, predicated instructions indicate whether the inactive destination vector elements are to undergo zeroing predication or merging predication. The type of predication is indicated by use of a qualifier suffix to the Governing predicate:
 - *Pg/Z* - zeroing predication.
 - *Pg/M* - merging predication.

Some instructions identify *Active elements* and *Inactive elements*, but do not write to a destination vector register. For these instructions, the Governing predicate operand is used with no zeroing or merging qualifier.

- Many *SVE* instructions have destructive instruction encodings. To avoid ambiguity, a constructive notation is used by the assembler language for these instructions. The destination register is repeated in the appropriate source operand position.

- SVE load/store addresses have a syntax that allows vector register operands within the address specified.
- A set of SVE aliases is defined for the AArch64 condition codes.

See also:

- *Structure of the A64 assembler language* in the ARM[®] *Architecture Reference Manual, ARMv8-A, for ARMv8-A architecture profile*
- [2.2 Process state, PSTATE N, Z, C and V Condition flags.](#)
- [11.9 Load/store/prefetch instructions.](#)

Draft

6.2 SVE ISA functional groups

I_{NSZQX} This section describes the functional groups for SVE instructions.

I_{KXTHC} SVE adds a set of instructions to the existing ARMv8-A A64 instruction set. The SVE instructions break down into the following functional groups:

- Load, store, and prefetch instructions.
- Integer operations.
- Vector address calculation.
- Bitwise operations.
- Floating-point operations.
- Predicate operations.
- Move operations.
- Reduction operations.

See also:

- *The A64 instruction set in the ARM[®] Architecture Reference Manual, ARMv8-A, for ARMv8-A architecture profile*

6.2.1 Load, store, and prefetch instructions

I_{QZSSG} SVE vector load and store instructions transfer data in memory to or from elements of one or more vector or predicate transfer registers. SVE also includes vector prefetch instructions that provide read and write hints to the memory system. For SVE predicated load, store, and prefetch instructions, the memory element access size and type that is associated with each vector element is specified by a suffix to the instruction mnemonic, independently of the element size of the transfer registers. For example, LD1SH. The following table shows the supported instruction suffixes for SVE load, store, and prefetch instructions:

Instruction suffix	Memory element access size and type
B	Unsigned byte
H	Unsigned halfword or half-precision floating-point
W	Unsigned word or single-precision floating-point
D	Unsigned doubleword or double-precision floating-point
SB	Signed byte
SH	Signed halfword
SW	Signed word

The element size of the transfer registers is always greater than or equal to the memory element access size. When the element size of the transfer registers is strictly greater than the memory element access size, then these are referred to as unpacked data accesses. In the case of unpacked data accesses:

- For load instructions, each element access is sign-extended or zero-extended to fill the vector element, according to its memory element access size and type.
- For store instructions, each vector element is truncated to the memory element access size.

Where the vector element size and the memory element access size are the same, then these are referred to as packed data accesses. Signed access types are not supported for packed data accesses. Packed and unpacked access sizes and types relate to the vector element size of the transfer registers, as defined in the following table:

Vector element	Packed access suffix	Unpacked access suffixes
.B	B	-
.H	H	B, SB
.S	W	H, SH, B, SB
.D	D	W, SW, H, SH, B, SB

For gather-load and scatter-store instructions, the vector element size can only be .S or .D. This means that any non-contiguous memory element access of less than a word is unpacked. Non-contiguous memory element accesses of a word can be either packed or unpacked, depending on the vector element size.

Load, store, and prefetch instructions consist of the following:

- 6.2.1.1 *Predicated single vector contiguous element accesses*
- 6.2.1.4 *Predicated replicating element loads*
- 6.2.1.2 *Predicated multiple vector contiguous structure load/store*
- 6.2.1.3 *Predicated non-contiguous element accesses*
- 6.2.1.4 *Predicated replicating element loads*
- 6.2.1.5 *Unpredicated vector register load/store*
- 6.2.1.6 *Unpredicated predicate register load/store*

All predicated load instructions zero the *Inactive elements* of the destination vector, except for Non-fault loads and First-fault loads when the corresponding FFR element is FALSE.

Prefetch instructions provide hints to hardware and do not change architectural state. Therefore, a Governing predicate for a prefetch instruction provides an additional hint which indicates the memory locations to be prefetched. Prefetch instructions require a prefetch operation specifier. SVE prefetch instructions support all of the prefetch operations except for the PLI prefetch operand types.

Load, store, and prefetch instructions that multiply a scalar index register or an index vector element by the memory element access size specify a shift type, followed by a shift amount in bits. The shift type can be one of LSL, SXTW, or UXTW. The shift amount is always Log2 of the memory element access size, in bytes. The shift amount defaults to zero when the memory element access size is a byte, and the shift size can be omitted. The shift type of LSL must be omitted if the shift amount is omitted.

When included as part of the assembler syntax for an instruction, MUL VL indicates that the specified immediate index value is multiplied by the size of the addressed vector or predicate in memory, measured in bytes, irrespective of predication. For a detailed description of the meaning of this assembler syntax for each instruction, see the appropriate subsection below.

When used in pseudocode, the symbol VL represents the vector length, measured in bits.

SVE load, store, and prefetch instructions do not support pre-indexed or post-indexed addressing.

6.2.1.1 Predicated single vector contiguous element accesses

‡_{SKT}WK

Predicated contiguous load and store instructions access memory locations starting from an address that is defined by a scalar base register plus either:

- A scalar index register.
- An immediate index value that is in the range -8 to 7, inclusive. This defaults to zero if omitted.

Predicated contiguous prefetch instructions address memory locations in a similar manner, with the index being either:

- A scalar index register.
- An immediate index value that is in the range of -32 to 31, inclusive. This defaults to zero if omitted.

For this group of SVE instructions:

- The immediate index value is a vector index, not an element index. The immediate index value is multiplied by the number of vector elements, irrespective of predication, and then multiplied by the memory element access size in bytes. The resulting offset is incremented following each element access by the memory element access size.
- The scalar index register value is multiplied by the memory element access size in bytes. The index value is incremented by one after each element access, but the scalar index register is not updated by the instruction.
- LD1* and ST1* instructions support both packed and unpacked data accesses, with a scalar index register or an immediate index value.
- First-fault load instructions that have the LDFF1* mnemonic prefix support both packed and unpacked data accesses, with a scalar index register that defaults to XZR if omitted.
- Non-fault load instructions that have the LDNF1* mnemonic prefix support both packed and unpacked data accesses, with an immediate index value.
- Non-temporal load instructions that have the LDNT1* mnemonic prefix and store instructions that have the STNT1* mnemonic prefix support only packed data accesses, with a scalar index register or an immediate index value.
- Prefetch instructions that have the PRF* mnemonic prefix support only packed data accesses, with a scalar index register or an immediate index value.
- When alignment checking is enabled for loads and stores, the value of the base address register must be aligned to the memory element access size.

Supported addressing modes	Assembler syntax
Scalar base + scalar index	[<Xn SP>, <Xm>{, LSL #<sh>}]
Scalar base + immediate index	[<Xn SP>{, #<simm>, MUL VL}]

6.2.1.2 Predicated multiple vector contiguous structure load/store

‡_{NSBCW}

Structure load instructions that have the LD2*, LD3*, or LD4* mnemonic prefix read N consecutive memory locations to the same-numbered element in each of the N vector transfer registers, where N = 2, 3, or 4, respectively. Structure store instructions that have the ST2*, ST3*, or ST4* mnemonic prefix write from the same-numbered element in each of the N consecutive vector transfer registers to N consecutive memory locations. The starting address is defined by a scalar base register plus either:

- A scalar index register.
- An immediate index that is a multiple of N, in the range $-8 \times N$ to $7 \times N$, inclusive. This defaults to zero if omitted.

For this group of SVE instructions:

- The immediate index value is a vector index, not an element index. The immediate index value is multiplied by the number of vector elements, irrespective of predication, and then multiplied by the memory element access size in bytes. The resulting offset is incremented following each element access by the memory element access size.
- The scalar index register value is multiplied by the memory element access size in bytes. Following each element access, the index value is incremented by one but the instruction does not update the scalar index register.
- Each predicate element applies to a single structure in memory, or equivalently to the same element number within each of the two, three, or four transferred vector registers.
- These instructions support packed data accesses only.
- When alignment checking is enabled for loads and stores, the base address must be aligned to the element access size.

Supported addressing modes	Assembler syntax
Scalar base + scalar index	[<XnSP>, <Xm>{, LSL #<sh>}]
Scalar base + immediate index	[<XnSP>{, #<simm>, MUL VL}]

6.2.1.3 Predicated non-contiguous element accesses

I_{RJWJ}

Predicated non-contiguous element accesses address non-contiguous memory locations that are specified by either:

- A scalar base register plus a vector of indices or offsets.
- A vector of base addresses plus an immediate byte offset. The immediate byte offset is a multiple of the memory element access size, in the range 0 to 31 times the memory element access size, inclusive, and defaults to zero if omitted.

For this group of *SVE* instructions:

- Vector registers used as part of the address must specify a vector element size of 32 bits or 64 bits, *.S* or *.D*. For load and store instructions, the transfer register must specify the same vector element size.
- If the index vector register contains 32-bit index values then the lowest 32 bits of each index vector element can either be zero-extended or sign-extended to 64 bits.
- For load and store instructions, the index vector elements are then optionally multiplied by the memory element access size, in bytes, if a shift amount is specified. For prefetch instructions the index vector elements are always multiplied by the memory element access size, in bytes.
- Non-contiguous LD1* instructions, ST1* instructions, and First-fault LDFF1* instructions support packed and unpacked data accesses. PRF* instructions only specify the memory element access size.
- When alignment checking is enabled for loads and stores, the computed virtual address of each element must be aligned to the memory element access size.

Supported addressing modes	Assembler syntax, 64-bit elements	Assembler syntax, 32-bit elements
Scalar base + 64-bit vector index	[<XnSP>, <Zm>.D{, LSL #<sh>}]	-
Scalar base + 32-bit vector index	[<XnSP>, <Zm>.D, (SIU)XTW { #<sh>}]	[<XnSP>, <Zm>.S, (SIU)XTW { #<sh>}]
Vector base + immediate offset	[<Zn>.D{, #<uimm>}]	[<Zn>.S{, #<uimm>}]

6.2.1.4 Predicated replicating element loads

I_{LQRQB}

The load and replicate instructions read one or more contiguous memory locations starting from an address that is defined by a scalar base register plus either:

- A scalar index register.
- An immediate byte offset.

This defaults to zero if omitted.

For this group of *SVE* instructions:

- The single element load and replicate, LD1R, *instructions load a single element value and replicate it into all Active elements** of the destination vector. These instructions support packed and unpacked data accesses. These instructions use an immediate byte offset that is a multiple of the memory element access size, in the range 0 to 63 times the memory element access size, inclusive.

- The 128-bit quadword load and replicate, LD1RQ*, instructions load a predicated 128-bit quadword vector segment from contiguous element values and replicate that segment into all segments of the destination vector. These instructions support only packed data accesses. These instructions can use a scalar index register that is multiplied by the memory element access size, or an immediate byte offset that is a multiple of 16, in the range of -128 to 112, inclusive.
- When alignment checking is enabled for loads and stores, the base address must be aligned to the memory element access size.

Supported addressing modes	Assembler syntax
Scalar base + scalar index	[<Xn SP>, <Xm>{, LSL #<sh>}]
Scalar base + immediate offset	[<Xn SP>{, #<imm>}]

6.2.1.5 Unpredicated vector register load/store

I_{YZYNG}

The unpredicated vector register load, LDR, and store, STR, instructions transfer a single vector register from or to memory locations that are specified by a scalar base register plus an immediate index value that is in the range -256 to 255, inclusive. The immediate index value defaults to zero if omitted. For this group of SVE instructions:

- The immediate index value is a vector index, not an element index. The immediate index value is multiplied by the current vector register length in bytes.
- The data transfer is performed as a contiguous stream of byte accesses in ascending element order, without endianness conversion.
- When alignment checking is enabled for loads and stores, the base address must be 16-byte aligned.

Supported addressing modes	Assembler syntax
Scalar base + immediate index	[<Xn SP>{, #<simm>, MUL VL}]

6.2.1.6 Unpredicated predicate register load/store

I_{PXNZC}

The unpredicated predicate register load, LDR, and store, STR, instructions transfer a single predicate register from or to memory locations that are specified by a scalar base register plus an immediate index value that is in the range -256 to 255, inclusive. The immediate index value defaults to zero if omitted. For this group of SVE instructions:

- The immediate index value is a predicate index, not an element index. The immediate index value is multiplied by the current predicate register length, in bytes.
- The data transfer is performed as a contiguous stream of byte accesses, each byte containing 8 consecutive predicate bits, in ascending bit and element order, without endian conversion.
- When alignment checking is enabled for loads and stores, the base address must be 2-byte aligned.

Supported addressing modes	Assembler syntax
Scalar base + immediate index	[<Xn SP>{, #<simm>, MUL VL}]

6.2.2 Vector move operations

6.2.2.1 Element move and broadcast

I_JSBCJ These instructions copy data from scalar registers, immediate values, and other vectors to selected vector elements. The copied data might be in an integer or floating-point format.

Mnemonic	Instruction	See
CPY	Copy signed integer immediate to vector elements	CPY (immediate)
	Copy general-purpose register to vector elements	CPY (scalar)
	Copy SIMD&FP scalar register to vector elements	CPY (SIMD&FP scalar)
DUP	Broadcast signed immediate to vector elements	DUP (immediate)
	Broadcast general-purpose register to vector elements	DUP (scalar)
FCPY	Copy 8-bit floating-point immediate to vector elements	FCPY
FDUP	Broadcast 8-bit floating-point immediate to vector elements	FDUP
FMOV	Move floating-point +0.0 to vector elements (unpredicated)	FMOV (zero, unpredicated)
	Move floating-point +0.0 to vector elements (predicated)	FMOV (zero, predicated)
	Move 8-bit floating-point immediate to vector elements (unpredicated)	FMOV (immediate, unpredicated)
	Move 8-bit floating-point immediate to vector elements (predicated)	FMOV (immediate, predicated)
MOV	Move signed integer immediate to vector elements (unpredicated)	MOV (immediate, unpredicated)
	Move signed integer immediate to vector elements (predicated)	MOV (immediate, predicated)
	Move general-purpose register to vector elements (unpredicated)	MOV (scalar, unpredicated)
	Move general-purpose register to vector elements (predicated)	MOV (scalar, predicated)
	Move SIMD&FP scalar register to vector elements (unpredicated)	MOV (SIMD&FP scalar, unpredicated)
	Move SIMD&FP scalar register to vector elements (predicated)	MOV (SIMD&FP scalar, predicated)
	Move vector register (unpredicated)	MOV (scalar, unpredicated)
Move vector register (predicated)	MOV (vector, predicated)	
SEL	Select vector elements from two vectors	SEL (vectors)

6.2.3 Integer operations

I_JZWFQ The following instructions operate on signed or unsigned integer data within a vector.

6.2.3.1 Integer arithmetic

I_FQSMH For binary operations, these instructions perform arithmetic operations on a source vector containing integer element values, and a second source vector of either integer element values or an immediate value. For ternary operations, these instructions perform arithmetic operations on a source vector containing integer element values, a

second source vector of either integer element values or an immediate value, and a third source vector containing integer element values.

Mnemonic	Instruction	See
ABS	Absolute value	ABS
ADD	Add vectors (predicated)	ADD (vectors, predicated)
	Add vectors (unpredicated)	ADD (vectors, unpredicated)
	Add immediate	ADD (immediate)
CNOT	Logically invert Boolean condition	CNOT
MAD	Multiply-add, writing to the multiplicand register	MAD
MLA	Multiply-add, writing to the addend register	MLA (vectors)
MLS	Multiply-subtract, writing to the addend register	MLS (vectors)
MSB	Multiply-subtract, writing to the multiplicand register	MSB
MUL	Multiply by immediate	MUL (immediate)
	Multiply vectors	MUL (vectors, predicated)
NEG	Negate	NEG
SABD	Signed absolute difference	SABD
SDIV	Signed divide	SDIV
SDIVR	Signed reverse divide	SDIVR
SMAX	Signed maximum with immediate	SMAX (immediate)
	Signed maximum vectors	SMAX (vectors)
SMIN	Signed minimum with immediate	SMIN (immediate)
	Signed minimum vectors	SMIN (vectors)
SMULH	Signed multiply returning high half	SMULH (predicated)
SQADD	Signed saturating add immediate	SQADD (immediate)
	Signed saturating add vectors	SQADD (vectors, unpredicated)
SQSUB	Signed saturating subtract immediate	SQSUB (immediate)
	Signed saturating subtract vectors	SQSUB (vectors, unpredicated)
SUB	Subtract immediate	SUB (immediate)
	Subtract vectors (predicated)	SUB (vectors, predicated)
	Subtract vectors (unpredicated)	SUB (vectors, unpredicated)
SUBR	Reversed subtract from immediate	SUBR (immediate)
	Reversed subtract vectors	SUBR (vectors)
SXTB	Signed byte extend	SXTB, SXTH, SXTW
SXTH	Signed halfword extend	SXTB, SXTH, SXTW
SXTW	Signed word extend	SXTB, SXTH, SXTW

Mnemonic	Instruction	See
UABD	Unsigned absolute difference	UABD
UDIV	Unsigned divide	UDIV
UDIVR	Unsigned reversed divide	UDIVR
UMAX	Unsigned maximum with immediate	UMAX (immediate)
	Unsigned maximum vectors	UMAX (vectors)
UMIN	Unsigned minimum with immediate	UMIN (immediate)
	Unsigned minimum vectors	UMIN (vectors)
UMULH	Unsigned multiply returning high half	UMULH (predicated)
UQADD	Unsigned saturating add immediate	UQADD (immediate)
	Unsigned saturating add vectors	UQADD (vectors, unpredicated)
UQSUB	Unsigned saturating subtract immediate	UQSUB (immediate)
	Unsigned saturating subtract vectors	UQSUB (vectors, unpredicated)
UXTB	Unsigned byte extend	UXTB, UXTH, UXTW
UXTH	Unsigned halfword extend	UXTB, UXTH, UXTW
UXTW	Unsigned word extend	UXTB, UXTH, UXTW

6.2.3.2 Integer dot product

I_{VRLQK}

The integer partial dot product instructions delimit the source vectors into groups of four 8-bit or 16-bit integer elements. Within each group of four elements, the elements in the first source vector are multiplied by the corresponding elements in the second source vector. The resulting widened products are summed and added to the 32-bit or 64-bit element of the accumulator and destination vector that aligns with the group of four elements in the first source vector.

The indexed forms of these instructions specify a single, numbered, group of four elements within each 128-bit segment of the second source vector as the multiplier for all the groups of four elements within the corresponding 128-bit segment of the first source vector.

The SUDOT and USDOT instructions are only supported ID_AA64ZFR0_EL1.I8MM is 1. The SUDOT and UDOT instructions only support groups of 8-bit elements.

Mnemonic	Instruction	See
SDOT	Signed dot product by vector	SDOT (vectors)
	Signed dot product by indexed elements	SDOT (indexed)
SUDOT	Signed by unsigned integer dot product by indexed elements	SUDOT
UDOT	Unsigned dot product by vector	UDOT (vectors)
	Unsigned dot product by indexed elements	UDOT (indexed)
USDOT	Unsigned by signed integer dot product	USDOT (vectors)

Mnemonic	Instruction	See
	Unsigned by signed integer dot product by indexed elements	USDOT (indexed)

6.2.3.3 Integer matrix multiply operations

I_{ZTKYV}

These instructions facilitate matrix multiplication and include integer matrix multiply-accumulate instructions.

The matrix multiplication instructions that are supported depend on the value of ID_AA64ZFR0_EL1.I8MM. The following table displays the supported instructions if ID_AA64ZFR0_EL1.I8MM is 1.

The matrix multiply-accumulate instructions delimit source and destination vectors into segments. Within each segment:

- The first source vector matrix is organized in row-by-row order.
- The second source vector matrix is organized in a column-by-column order.
- The destination vector matrix is organized in row-by-row order.

One matrix multiplication is performed per vector segment. For other matrix multiply instructions, the vector segment length is 128 bits.

Mnemonic	Instruction	See
SMMLA	Widening signed 8-bit integer matrix multiply-accumulate into 2x2 matrix	SMMLA
UMMLA	Widening unsigned 8-bit integer matrix multiply-accumulate into 2x2 matrix	UMMLA
USMMLA	Widening mixed sign 8-bit integer matrix multiply-accumulate into 2x2 matrix	USMMLA

6.2.3.4 Integer comparisons

I_{ZTCWQ}

These instructions compare *Active elements* in the first source vector with the corresponding elements in a second vector or with an immediate value. The Boolean result of each comparison is placed in the corresponding element of the destination predicate. *Inactive elements* in the destination predicate register are set to zero. All integer comparisons set the N, Z, and C condition flags based on the predicate result, and set the V flag to zero. No floating-point comparisons modify any condition flags.

The wide element variants of the compare instructions allow a packed vector of narrower elements to be compared with wider 64-bit elements. These instructions treat the second source vector as having a fixed 64-bit doubleword element size and compare each narrow element of the first source vector with the corresponding vertically-aligned wide element of the second source vector. For example, if the first source vector contained 8-bit byte elements, then 8-bit element[0] to element[7] of the first source vector are compared with 64-bit element[0] of the second source vector, 8-bit element[8] to element[15] with 64-bit element[1], and so on. All 64 bits of the wide elements are significant for the comparison, with the narrow elements being sign-extended or zero-extended to 64 bits as appropriate for the type of comparison.

Mnemonic	Instruction	See
CMPEQ	Compare signed equal to immediate	CMP<cc> (immediate)
	Compare signed equal to wide elements	CMP<cc> (wide elements)
	Compare signed equal to vector	CMP<cc> (vectors)

Mnemonic	Instruction	See
CMPGE	Compare signed greater than or equal to immediate	CMP<cc> (immediate)
	Compare signed greater than or equal to wide elements	CMP<cc> (wide elements)
	Compare signed greater than or equal to vector	CMP<cc> (vectors)
CMPGT	Compare signed greater than immediate	CMP<cc> (immediate)
	Compare signed greater than wide elements	CMP<cc> (wide elements)
	Compare signed greater than vector	CMP<cc> (vectors)
CMPHI	Compare unsigned higher than immediate	CMP<cc> (immediate)
	Compare unsigned higher than wide elements	CMP<cc> (wide elements)
	Compare unsigned higher than vector	CMP<cc> (vectors)
CMPHS	Compare unsigned higher than or same as immediate	CMP<cc> (immediate)
	Compare unsigned higher than or same as wide elements	CMP<cc> (wide elements)
	Compare unsigned higher than or same as vector	CMP<cc> (vectors)
CMPLE	Compare signed less than or equal to immediate	CMP<cc> (immediate)
	Compare signed less than or equal to wide elements	CMP<cc> (wide elements)
	Compare signed less than or equal to vector	CMP<cc> (vectors)
CMPLO	Compare unsigned lower than immediate	CMP<cc> (immediate)
	Compare unsigned lower than 64-bit wide elements	CMP<cc> (wide elements)
	Compare unsigned lower than vector	CMP<cc> (vectors)
CMPLS	Compare unsigned lower or same as immediate	CMP<cc> (immediate)
	Compare unsigned lower or same as wide elements	CMP<cc> (wide elements)
	Compare unsigned lower or same as vector	CMP<cc> (vectors)
CMPLT	Compare signed less than immediate	CMP<cc> (immediate)
	Compare signed less than wide elements	CMP<cc> (wide elements)
	Compare signed less than vector	CMP<cc> (vectors)
CMPNE	Compare not equal to immediate	CMP<cc> (immediate)
	Compare not equal to wide elements	CMP<cc> (wide elements)
	Compare not equal to vector	CMP<cc> (vectors)

6.2.3.5 Vector address calculation

I_{SSWB}T

These instructions compute vectors of addresses and addresses of vectors. This includes instructions to add a multiple of the current vector length or predicate register length, in bytes, to a general-purpose register.

The ADR instruction is an integer arithmetic operation that is used to calculate a vector of 64-bit or 32-bit addresses.

The destination register elements are computed by the addition of the corresponding elements in the source registers, with an optional sign or zero extension and optional bitwise left shift of 1-3 bits applied to the final operands. This can be considered as the addition of a vector base and a scaled vector index.

32-bit addresses are computed by the addition of a 32-bit base and a scaled 32-bit unsigned index.

64-bit addresses are computed by one of:

- Addition of a 64-bit base and a scaled 64-bit unsigned index.
- Addition of a 64-bit base and a scaled, zero-extended 32-bit index.
- Addition of a 64-bit base and a scaled, sign-extended 32-bit index.

Mnemonic	Instruction	See
ADDVL	Add multiple of vector length, in bytes, to scalar register	ADDVL
ADDPL	Add multiple of predicate register length, in bytes, to scalar register	ADDPL
ADR	Compute vector of addresses	ADR
RDVL	Read multiple of vector register length, in bytes, to scalar register	RDVL

6.2.4 Bitwise operations

6.2.4.1 Bitwise logical operations

I_{JTPLZ} These instructions perform bitwise logical operations on vectors. Where operations are unpredicated, the operations are independent of the element size.

Mnemonic	Instruction	See
AND	Bitwise AND vectors (predicated)	AND (vectors, predicated)
	Bitwise AND vectors (unpredicated)	AND (vectors, unpredicated)
	Bitwise AND with immediate	AND (immediate)
BIC	Bitwise clear with vector (predicated)	BIC (vectors, predicated)
	Bitwise clear with vector (unpredicated)	BIC (vectors, unpredicated)
	Bitwise clear using immediate	BIC (immediate)
DUPM	Broadcast bitmask immediate to vector (unpredicated)	DUPM
EON	Bitwise exclusive OR with inverted immediate	EON
EOR	Bitwise exclusive OR vectors (predicated)	EOR (vectors, predicated)
	Bitwise exclusive OR vectors (unpredicated)	EOR (vectors, unpredicated)
	Bitwise exclusive OR with immediate	EOR (immediate)
MOV	Move bitmask immediate to vector	MOV (bitmask immediate)
	Move vector register	MOV (vector, unpredicated)
NOT	Bitwise invert vector	NOT (vector)
ORN	Bitwise OR with inverted immediate	ORN (immediate)
ORR	Bitwise OR vectors (predicated)	ORR (vectors, predicated)
	Bitwise OR vectors (unpredicated)	ORR (vectors, unpredicated)
	Bitwise OR with immediate	ORR (immediate)

6.2.4.2 Bitwise shift, reverse, and count

‡_{ZMLFJ}

Bitwise shifts, reversals, and counts within vector elements.

Shift counts saturate at the number of bits per element, rather than being used modulo the element size. If modulo behavior is required, then the modulus must be computed separately.

The wide element variants of the bitwise shift instructions allow a packed vector of narrower elements to be shifted by wider 64-bit shift amounts. These instructions treat the second source vector as having a fixed 64-bit doubleword element size and shift each narrow element of the first source vector by the corresponding vertically-aligned wide element of the second source vector. For example, if the first source vector contained 8-bit byte elements, then 8-bit element[0] to element[7] of the first vector are shifted by 64-bit element[0] of the second source vector, 8-bit element [8] to element[15] by 64-bit element[1], and so on. All 64 bits of the wide shift amount are significant.

Mnemonic	Instruction	See
ASR	Arithmetic shift right by immediate (predicated)	ASR (immediate, predicated)
	Arithmetic shift right by immediate (unpredicated)	ASR (immediate, unpredicated)
	Arithmetic shift right by wide elements (predicated)	ASR (wide elements, predicated)
	Arithmetic shift right by wide elements (unpredicated)	ASR (wide elements, unpredicated)
	Arithmetic shift right by vector	ASR (immediate, predicated)
ASRD	Arithmetic shift right for divide by immediate	ASRD
ASRR	Reversed arithmetic shift right by vector	ASRR
CLS	Count leading sign bits	CLS
CLZ	Count leading zero bits	CLZ
CNT	Count nonzero bits	CNT
LSL	Logical shift left by immediate (predicated)	LSL (immediate, predicated)
	Logical shift left by immediate (unpredicated)	LSL (immediate, unpredicated)
	Logical shift left by wide elements (predicated)	LSL (wide elements, predicated)
	Logical shift left by wide elements (unpredicated)	LSL (wide elements, unpredicated)
	Logical shift left by vector	LSL (vectors)
LSLR	Reversed logical shift left by vector	LSLR
LSR	Logical shift right by immediate (predicated)	LSR
	Logical shift right by immediate (unpredicated)	LSR (immediate, unpredicated)
	Logical shift right by wide elements (predicated)	LSR (wide elements, predicated)
	Logical shift right by wide elements (unpredicated)	LSR (wide elements, unpredicated)
	Logical shift right by vector	LSR (vectors)
LSRR	Reversed logical shift right by vector	LSRR
RBIT	Reverse bits	RBIT

6.2.5 Floating-point operations

\perp_{JXJHX} The following instructions operate on floating-point data within a vector.

6.2.5.1 BFloat16 floating-point support

\perp_{LZCLK} BFloat16, or BF16, is a 16-bit floating-point storage format defined by the Arm architecture such that it inherits many of its properties and behaviors from the IEEE 754 single-precision format. For more information, see the section titled BFloat16 floating-point format in the Arm® Architecture Reference Manual, Armv8-A, for Armv8-A architecture profile.

The BFloat16 instructions are only supported if ID_AA64ZFR0_EL1.BF16 is not 0.

6.2.5.2 Floating-point arithmetic

\perp_{SVWLR} These instructions perform arithmetic operations on vectors containing floating-point element values.

Mnemonic	Instruction	See
FABD	Floating-point absolute difference	FABD
FABS	Floating-point absolute value	FABS
FADD	Floating-point add immediate	FADD (immediate)
	Floating-point add (predicated)	FADD (vectors, predicated)
	Floating-point add (unpredicated)	FADD (vectors, unpredicated)
FDIV	Floating-point divide	FDIV
FDIVR	Floating-point reversed divide	FDIVR
FMAX	Floating-point maximum with immediate	FMAX (immediate)
	Floating-point maximum vectors	FMAX (vectors)
FMAXNM	Floating-point maximum number with immediate	FMAXNM (immediate)
	Floating-point maximum number vectors	FMAXNM (vectors)
FMIN	Floating-point minimum with immediate	FMIN (immediate)
	Floating-point minimum vectors	FMIN(vectors)
FMINNM	Floating-point minimum number with immediate	FMINNM (immediate)
	Floating-point minimum number vectors	FMINNM (vectors)
FMUL	Floating-point multiply by immediate	FMUL (immediate)
	Floating-point multiply vectors (predicated)	FMUL (vectors, predicated)
	Floating-point multiply vectors (unpredicated)	FMUL (vectors, unpredicated)
FMULX	Floating-point multiply-extended	FMULX
FNEG	Floating-point negate	FNEG
FRECPE	Floating-point reciprocal estimate	FRECPE
FRECPS	Floating-point reciprocal step	FRECPS
FRECPX	Floating-point reciprocal exponent	FRECPX
FRSQRTE	Floating-point reciprocal square root estimate	FRSQRTE
FRSQRTS	Floating-point reciprocal square root step	FRSQRTS

Mnemonic	Instruction	See
FSCALE	Floating-point adjust exponent by vector	FSCALE
FSQRT	Floating-point square root	FSQRT
FSUB	Floating-point subtract immediate	FSUB (immediate)
	Floating-point subtract vectors (predicated)	FSUB (vectors, predicated)
	Floating-point subtract vectors (unpredicated)	FSUB (vectors, unpredicated)
FSUBR	Floating-point reversed subtract from immediate	FSUBR (immediate)
	Floating-point reversed subtract vectors	FSUBR(vectors)

6.2.5.3 Floating-point multiply accumulate

I_{DWMYJ}

These instructions perform floating-point fused multiply-add or multiply-subtract operations and their negated forms. There are two groups of these instructions, as follows:

- Instructions where the result of the operation is written to the addend register.
 - Supported instructions are: FMLA, FMLS, FNMLA, FNMLS.
- Instructions where the result of the operation is written to the multiplicand register.
 - Supported instructions are: FMAD, FMSB, FNMAD, FNMSB.

Mnemonic	Instruction	See
FMLA	Floating-point fused multiply-add vectors, writing to the addend	FMLA (vectors)
FMLS	Floating-point fused multiply-subtract vectors, writing to the addend	FMLS (vectors)
FNMLA	Floating-point negated fused multiply-add vectors, writing to the addend	FNMLA
FNMLS	Floating-point negated fused multiply-subtract vectors, writing to the addend	FNMLS
FMAD	Floating-point fused multiply-add vectors, writing to the multiplicand	FMAD
FMSB	Floating-point fused multiply-subtract vectors, writing to the multiplicand	FMSB
FNMAD	Floating-point negated fused multiply-add vectors, writing to the multiplicand	FNMAD
FNMSB	Floating-point negated fused multiply-subtract vectors, writing to the multiplicand	FNMSB

6.2.5.4 Floating-point complex arithmetic

I_{MGRHZ}

These instructions perform arithmetic on vectors containing floating-point complex numbers as interleaved pairs of elements, where the even-numbered elements contain the real components and the odd-numbered elements contain the imaginary components.

The FCADD instructions rotate the complex numbers in the second source vector by 90 degrees or 270 degrees in the direction from the positive real axis towards the positive imaginary axis, when considered in polar representation, before adding active pairs of elements to the corresponding elements of the first source vector in a destructive manner.

The FCMLA instructions perform a transformation of the operands to allow the creation of multiply-add or multiply-subtract operations on complex numbers by combining two of the instructions. The transformations performed are as follows:

- The complex numbers in the second source vector, considered in polar form, are rotated by 0 degrees or 180 degrees before multiplying by the duplicated real components of the first source vector.
- The complex numbers in the second source vector, considered in polar form, are rotated by 90 degrees or 270 degrees before multiplying by the duplicated imaginary components of the first source vector.

The resulting products are then added to the corresponding components of the destination and addend vector, without intermediate rounding. Two FCMLA instructions can be used as follows:

FCMLA Zda.S, Pg/M, Zn.S, Zm.S, #A . . . FCMLA Zda.S, Pg/M, Zn.S, Zm.S, #B

For example, some meaningful combinations of A and B are:

- A=0, B=90. In this case, the two vectors of complex numbers in Zn and Zm are multiplied and the products are added to the complex numbers in Zda.
- A=0, B=270. In this case, the conjugates of the complex numbers in Zn are multiplied by the complex numbers in Zm and the products are added to the complex numbers in Zda.
- A=180, B=270. In this case, the two vectors of complex numbers in Zn and Zm are multiplied and the products are subtracted from the complex numbers in Zda.
- A=180, B=90. In this case, the conjugates of the complex numbers in Zn are multiplied by the complex numbers in Zm and the products are subtracted from the complex numbers in Zda.

The lack of intermediate rounding can give unexpected results in certain cases relative to a traditional sequence of independent multiply, add, and subtract instructions.

In addition, when using these instructions, the behavior of calculations such as $(\infty+\infty)$ multiplied by $(0+i)$ is $(NaN+NaNi)$, rather than the result expected by ISO C, which is complex ∞ . The expectation is that these instructions are only used in situations where the effect of differences in the rounding and handling of infinities are not material to the calculation.

Mnemonic	Instruction	See
FCADD	Floating-point complex add with rotate	FCADD
FCMLA	Floating-point complex multiply-add with rotate	FCMLA (vectors)

6.2.5.5 Floating-point rounding and conversion

I_{TMXZV}

These instructions change floating-point size and precision, round floating-point to integral floating-point with explicit rounding mode, and convert floating-point to or from integer format.

Mnemonic	Instruction	See
BFCVT	Floating-point down convert to BFloat16 format	BFCVT
BFCVTNT	Floating-point down convert and narrow to BFloat16 format (top, predicated)	BFCVTNT
FCVT	Floating-point convert precision	FCVT
FCVTZS	Floating-point convert to signed integer, rounding toward zero	FCVTZS
FCVTZU	Floating-point convert to unsigned integer, rounding toward zero	FCVTZU
FRINTA	Floating-point round to integral value, to nearest with ties away from zero	FRINT<r>
FRINTI	Floating-point round to integral value, using the current rounding mode	FRINT<r>

Mnemonic	Instruction	See
FRINTM	Floating-point round to integral value, toward minus infinity	FRINT<rd>
FRINTN	Floating-point round to integral value, to nearest with ties to even	FRINT<rd>
FRINTP	Floating-point round to integral value, toward plus infinity	FRINT<rd>
FRINTX	Floating-point round to integral value exact, using the current rounding mode	FRINT<rd>
FRINTZ	Floating-point round to integral value, toward zero	FRINT<rd>
SCVTF	Signed integer convert to floating-point	SCVTF
UCVTF	Unsigned integer convert to floating-point	UCVTF

6.2.5.6 Floating-point comparisons

I_{KBYDX}

These instructions compare active floating-point element values in the first source vector with corresponding elements in the second vector or with the immediate value +0.0. The Boolean result of each comparison is placed in the corresponding element of the destination predicate. *Inactive elements* in the destination predicate register are set to zero. Floating-point vector comparisons do not set the condition flags.

Mnemonic	Instruction	See
FACGE	Floating-point absolute compare greater than or equal	FAC<cc>
FACGT	Floating-point absolute compare greater than	FAC<cc>
FACLE	Floating-point absolute compare less than or equal	FACLE
FACLT	Floating-point absolute compare less than	FACLT
FCMEQ	Floating-point compare equal to zero	FCM<cc> (zero)
	Floating-point compare equal to vector	FCM<cc> (vectors)
FCMGE	Floating-point compare greater than or equal to zero	FCM<cc> (zero)
	Floating-point compare greater than or equal to vector	FCM<cc> (vectors)
FCMGT	Floating-point compare greater than zero	FCM<cc> (zero)
	Floating-point compare greater than vector	FCM<cc> (vectors)
FCMLE	Floating-point compare less than or equal to zero	FCM<cc> (zero)
	Floating-point compare less than or equal to vector	FCM<cc> (vectors)
FCMLT	Floating-point compare less than zero	FCM<cc> (zero)
	Floating-point compare less than vector	FCM<cc> (vectors)
FCMNE	Floating-point compare not equal to zero	FCM<cc> (zero)
	Floating-point compare not equal to vector	FCM<cc> (vectors)
FCMUO	Floating-point unordered vectors	FCM<cc> (vectors)

6.2.5.7 Floating-point transcendental acceleration

I_{LRHVT} The floating-point transcendental instructions accelerate calculations of sine, cosine, and exponential functions for vectors containing floating-point element values.

The trigonometric instructions accelerate the calculation of a polynomial series approximation for the sine and cosine functions. The exponential instruction accelerates the polynomial series calculation of the exponential function.

Mnemonic	Instruction	See
FTMAD	Floating-point trigonometric multiply-add coefficient	FTMAD
FTSMUL	Floating-point trigonometric starting value	FTSMUL
FTSSEL	Floating-point trigonometric select coefficient	FTSSEL
FEXPA	Floating-point exponential accelerator	FEXPA

6.2.5.8 Floating-point indexed multiplies

I_{ZYHHJ} These instructions multiply all floating-point elements within each 128-bit segment of the first source vector by the single numbered element within the corresponding segment of the second source vector. For the FMLA and FMLS instructions, the products are destructively added or subtracted from the corresponding elements of the addend and destination vector, without intermediate rounding.

Mnemonic	Instruction	See
FMLA	Floating-point fused multiply-add by indexed elements	FMLA (indexed)
FMLS	Floating-point fused multiply-subtract by indexed elements	FMLS (indexed)
FMUL	Floating-point multiply by indexed elements	FMUL (indexed)

6.2.5.9 Floating-point matrix multiply operations

I_{QSQWL} Instructions to facilitate matrix multiplication include floating-point matrix multiply-accumulate instructions, and companion instructions that support data rearrangements in vector registers as required by some of the matrix multiply instructions.

The floating-point matrix multiplication instructions and companion instructions that are supported depend on the values of:

- If FEAT_F64MM is implemented, ID_AA64ZFR0_EL1.F64MM, as shown in Instructions supported when ID_AA64ZFR0_EL1.F64MM is 1.
- If FEAT_F32MM is implemented, ID_AA64ZFR0_EL1.F32MM, as shown in Instructions supported when ID_AA64ZFR0_EL1.F32MM is 1.

The matrix multiply-accumulate instructions delimit source and destination vectors into segments. Within each segment:

- The first source vector matrix is organized in row-by-row order.
- The second source vector matrix is organized in a column-by-column order.
- The destination vector matrix is organized in row-by-row order.

One matrix multiplication is performed per vector segment. For the double-precision matrix multiply instructions, the vector segment length and minimum vector length is 256 bits. Double-precision matrix multiply instructions

are not supported when the vector length is 128. For other matrix multiply instructions, the vector segment length is 128 bits.

The floating-point matrix multiply-accumulate instructions perform all arithmetic with IEEE 754 compliant numerical behaviors and observe the FPCR controls. These are all detailed in the Arm® Architecture Reference Manual, Armv8-A, for Armv8-A architecture profile in the following sections:

- Single-precision floating-point format.
- Double-precision floating-point format.
- Advanced SIMD and floating-point support.
- FPCR, Floating-point Control Register.

For the floating-point matrix multiply-accumulate instructions, the order of accumulations is architecturally defined. The multiplications and additions are not fused, so intermediate rounding is performed after every multiplication and every addition.

†_{TJXGZ}

The following table shows the floating-point matrix multiplication instructions and companion instructions that are supported if ID_AA64ZFR0_EL1.F64MM is 1:

Mnemonic	Instruction	See
FMMLA	Floating-point matrix multiply-accumulate into 2x2 matrix (double-precision)	FMMLA
LD1ROB	Contiguous load and replicate thirty-two bytes, scalar plus scalar	LD1ROB (scalar plus scalar)
	Contiguous load and replicate thirty-two bytes, scalar plus immediate	LD1ROB (scalar plus immediate)
LD1ROD	Contiguous load and replicate four doublewords, scalar plus scalar	LD1ROD (scalar plus scalar)
	Contiguous load and replicate four doublewords, scalar plus immediate	LD1ROD (scalar plus immediate)
LD1ROH	Contiguous load and replicate sixteen halfwords, scalar plus scalar	LD1ROH (scalar plus scalar)
	Contiguous load and replicate sixteen halfwords, scalar plus immediate	LD1ROH (scalar plus immediate)
LD1ROW	Contiguous load and replicate eight words, scalar plus scalar	LD1ROW (scalar plus scalar)
	Contiguous load and replicate eight words, scalar plus immediate	LD1ROW (scalar plus immediate)
TRN1, TRN2	Interleave even or odd 128-bit elements from two vectors	TRN1, TRN2 (vectors)
UZP1, UZP2	Concatenate even or odd 128-bit elements from two vectors	UZP1, UZP2 (vectors)
ZIP1, ZIP2	Interleave 128-bit elements from two half vectors	ZIP1, ZIP2 (vectors)

†_{PTVQV}

The following table shows the floating-point matrix multiplication instructions and companion instructions that are supported if ID_AA64ZFR0_EL1.F32MM is 1:

Mnemonic	Instruction	See
FMMLA	Floating-point matrix multiply-accumulate into 2x2 matrix (double-precision)	FMMLA

6.2.5.10 BFloat16 floating-point instructions

I_{ZLRKW}

All of these instructions perform an implicit conversion of vectors of BF16 input values to IEEE 754 single-precision floating-point format. In addition, the BFDOT and BFMLLA instructions perform an N-way dot-product calculation that accumulates the products into a vector of single-precision accumulators.

All of these instructions perform arithmetic with fixed numeric behaviors. For more information, see the section titled BFloat16 floating-point format in the Arm® Architecture Reference Manual, Armv8-A, for Armv8-A architecture profile.

The BFloat16 instructions are only supported if ID_AA64ZFR0_EL1.BF16 is 1.

Mnemonic	Instruction	See
BFDOT	BFloat16 floating-point dot product by vector	BFDOT (vectors)
	BFloat16 floating-point dot product by indexed elements	BFDOT (indexed)
BFMLLA	BFloat16 floating-point matrix multiply-accumulate into 2x2 matrix	BFMLLA
BFMLALB	BFloat16 floating-point widening multiply accumulate long bottom by vector	BFMLALB (vectors)
	BFloat16 floating-point widening multiply accumulate long bottom by indexed elements	BFMLALB (indexed)
BFMLALT	BFloat16 floating-point widening multiply accumulate long top by vector	BFMLALT (vectors)
	BFloat16 floating-point widening multiply accumulate long bottom by indexed elements	BFMLALT (vectors)

6.2.6 Predicate operations

I_{LDPLK}

These instructions relate to operations that manipulate the predicate registers.

Some of these instructions are insensitive to the predicate element size and specify an explicit byte element size qualifier, .B, but an assembler must accept any qualifier, or none.

6.2.6.1 Predicate initialization

I_{RNWTM}

These instructions initialize predicate elements.

Predicate elements can be initialized to be FALSE, or to be TRUE when their element number is less than:

- A fixed number of elements from the following range: VL1-VL8, VL16, VL32, VL64, VL128 or VL256.
- The largest power of two elements, POW2.
- The largest multiple of three or four elements, MUL3 or MUL4.
- The number of accessible elements, ALL, which is implicitly a multiple of two.

Unspecified or out of range constraint encodings generate a predicate with values that are all FALSE and do not cause an Undefined Instruction exception.

Mnemonic	Instruction	See
PFALSE	Set all predicate elements to FALSE	PFALSE
PTRUE	Initialize predicate elements from named constraint	PTRUE, PTRUES
PTRUES	Initialize predicate elements from named constraint, setting the condition flags	PTRUE, PTRUES

6.2.6.2 Predicate move operations

IQFBXX

These instructions operate on all bits of the predicate registers, implying a fixed, 1-bit predicate element size. The flag-setting variants set the N, Z, and C condition flags based on the predicate result, and set the V flag to zero. Because these instructions operate with a fixed, 1-bit element size, the Governing predicate for the flag-setting variants should be in the canonical form for a predicate element size in order to generate a meaningful set of condition flags for that element size.

Mnemonic	Instruction	See
SEL	Select predicate elements from two predicates	SEL (predicates)
MOV	Move predicate elements (predicated, merging)	MOV (predicate, predicated, merging)
	Move predicate elements (predicated, zeroing)	MOV (predicate, predicated, zeroing)
	Move predicate elements (unpredicated)	MOV (predicate, unpredicated)
MOVS	Move predicate elements, setting the condition flags (predicated)	MOVS (predicated)
	Move predicate elements, setting the condition flags (unpredicated)	MOVS (unpredicated)

6.2.6.3 Predicate logical operations

IJGHMH

These instructions perform bitwise logical operations on predicate registers that operate on all bits of the register, implying a fixed, 1-bit predicate element size. The flag-setting variants set the N, Z, and C condition flags based on the predicate result, and set the V flag to zero. Inactive elements in the destination Predicate register are set to zero, except for PTEST which does not specify a destination register. Because these instructions operate with a fixed, 1-bit element size, the Governing predicate for the flag-setting variants should be in the canonical form for a predicate element size in order to generate a meaningful set of condition flags for that element size.

Mnemonic	Instruction	See
AND	Bitwise AND predicates	AND
ANDS	Bitwise AND predicates, setting the condition flags	ANDS
BIC	Bitwise clear predicates	BIC
BICS	Bitwise clear predicates, setting the condition flags	BICS
EOR	Bitwise exclusive OR predicates	EOR

Mnemonic	Instruction	See
EORS	Bitwise exclusive OR predicates, setting the condition flags	EORS
NAND	Bitwise NAND predicates	NAND
NANDS	Bitwise NAND predicates, setting the condition flags	NANDS
NOR	Bitwise NOR predicates	NOR
NORS	Bitwise NOR predicates, setting the condition flags	NORS
NOT	Bitwise invert predicate	NOT
NOTS	Bitwise invert predicate, setting the condition flags	NOTS
ORN	Bitwise OR inverted predicate	ORN
ORNS	Bitwise OR inverted predicate, setting the condition flags	ORNS
ORR	Bitwise OR predicates	ORR
ORRS	Bitwise OR predicates, setting the condition flags	ORRS
PTEST	Test predicate value, setting the condition flags	PTEST

6.2.6.4 FFR predicate handling

┆_{LTyHB}

These instructions work with *SVE* First-fault and Non-fault loads using the FFR to determine which elements have been successfully loaded and which remain to be loaded on a subsequent iteration. The RDIFFRS instruction sets the N, Z, and C condition flags based on the predicate result, and sets the V flag to zero. Because these instructions operate with a fixed, 1-bit element size, the Governing predicate for the RDIFFRS instruction should be in the canonical form for a predicate element size in order to generate a meaningful set of condition flags for that element size.

Mnemonic	Instruction	See
RDIFFR	Return predicate of successfully loaded elements (unpredicated)	RDIFFR
	Return predicate of successfully loaded elements (predicated)	RDIFFR
RDIFFRS	Return predicate of successfully loaded elements, setting the condition flags (predicated)	RDIFFRS
SETFFR	Initialize the First-fault register to all TRUE	SETFFR
WRFFR	Write a predicate register to the First-fault register	WRFFR

6.2.6.5 Predicate counts

┆_{DMNRB}

These instructions count either the number of Active predicate elements that are set to TRUE, or the number of elements implied by a named predicate constraint. The count can be placed in a general-purpose register, or used to increment or decrement a vector or general-purpose register.

Signed or unsigned saturating variants handle cases where, for example, an increment might cause a vectorized scalar loop index to overflow and therefore never satisfy a loop termination condition that compares it with a limit that is close to the maximum integer value.

The named predicate constraint limits the number of elements to:

- A fixed number of elements from the following range: VL1-VL8, VL16, VL32, VL64, VL128 or VL256.
- The largest power of two elements, POW2.
- The largest multiple of three or four elements, MUL3 or MUL4.
- The number of accessible elements, ALL, implicitly a multiple of two.

Unspecified or out of range predicate constraint encodings generate a zero element count and do not cause an Undefined Instruction exception.

Mnemonic	Instruction	See
CNTB	Set scalar to multiple of 8-bit predicate constraint element count	CNTB
CNTH	Set scalar to multiple of 16-bit predicate constraint element count	CNTH
CNTW	Set scalar to multiple of 32-bit predicate constraint element count	CNTW
CNTD	Set scalar to multiple of 64-bit predicate constraint element count	CNTD
CNTP	Set scalar to the number of Active predicate elements that are TRUE	CNTP
DECB	Decrement scalar by multiple of 8-bit predicate constraint element count	DECB
DECH	Decrement scalar by multiple of 16-bit predicate constraint element count	DECH (scalar)
	Decrement vector by multiple of 16-bit predicate constraint element count	DECH (vector)
DECW	Decrement scalar by multiple of 32-bit predicate constraint element count	DECW (scalar)
	Decrement vector by multiple of 32-bit predicate constraint element count	DECW (vector)
DECD	Decrement scalar by multiple of 64-bit predicate constraint element count	DECD (scalar)
	Decrement vector by multiple of 64-bit predicate constraint element count	DECD (vector)
DECP	Decrement scalar by the number of predicate elements that are TRUE	DECP (scalar)
	Decrement vector by the number of Active predicate elements that are TRUE	DECP (vector)
INCB	Increment scalar by multiple of 8-bit predicate constraint element count	INCB (scalar)
INCH	Increment scalar by multiple of 16-bit predicate constraint element count	INCH (scalar)
	Increment vector by multiple of 16-bit predicate constraint element count	INCH (vector)
INCW	Increment scalar by multiple of 32-bit predicate constraint element count	INCW (scalar)
	Increment vector by multiple of 32-bit predicate constraint element count	INCW (vector)
INCD	Increment scalar by multiple of 64-bit predicate constraint element count	INCD (scalar)
	Increment vector by multiple of 64-bit predicate constraint element count	INCD (vector)
INCP	Increment scalar by the number of predicate elements that are TRUE	INCP (scalar)
	Increment vector by the number of predicate elements that are TRUE	INCP (vector)
SQDECB	Signed saturating decrement scalar by multiple of 8-bit predicate constraint element count	SQDECB
SQDECH	Signed saturating decrement scalar by multiple of 16-bit predicate constraint element count	SQDECH (scalar)
	Signed saturating decrement vector by multiple of 16-bit predicate constraint element count	SQDECH (vector)
SQDECW	Signed saturating decrement scalar by multiple of 32-bit predicate constraint element count	SQDECW (scalar)

Mnemonic	Instruction	See
	Signed saturating decrement vector by multiple of 32-bit predicate constraint element count	SQDECW (vector)
SQDECD	Signed saturating decrement scalar by multiple of 64-bit predicate constraint element count	SQDECD (scalar)
	Signed saturating decrement vector by multiple of 64-bit predicate constraint element count	SQDECD (vector)
SQDECP	Signed saturating decrement scalar the number of predicate elements that are TRUE	SQDECP (scalar)
	Signed saturating decrement vector by the number of predicate elements that are TRUE	SQDECP (vector)
SQINCB	Signed saturating increment scalar by multiple of 8-bit predicate constraint element count	SQINCB (scalar)
SQINCH	Signed saturating increment scalar by multiple of 16-bit predicate constraint element count	SQINCH (scalar)
	Signed saturating increment vector by multiple of 16-bit predicate constraint element count	SQINCH (vector)
SQINCW	Signed saturating increment scalar by multiple of 32-bit predicate constraint element count	SQINCW (scalar)
	Signed saturating increment vector by multiple of 32-bit predicate constraint element count	SQINCW (vector)
SQINCD	Signed saturating increment scalar by multiple of 64-bit predicate constraint element count	SQINCD (scalar)
	Signed saturating increment vector by multiple of 64-bit predicate constraint element count	SQINCD (vector)
SQINCP	Signed saturating increment scalar by the number of predicate elements that are TRUE	SQINCP (scalar)
	Signed saturating increment vector by the number of predicate elements that are TRUE	SQINCP (vector)
UQDECB	Unsigned saturating decrement scalar by multiple of 8-bit predicate constraint element count	UQDECB
UQDECH	Unsigned saturating decrement scalar by multiple of 16-bit predicate constraint element count	UQDECH (scalar)
	Unsigned saturating decrement vector by multiple of 16-bit predicate constraint element count	UQDECH (vector)
UQDECW	Unsigned saturating decrement scalar by multiple of 32-bit predicate constraint element count	UQDECW (scalar)
	Unsigned saturating decrement vector by multiple of 32-bit predicate constraint element count	UQDECW (vector)
UQDECD	Unsigned saturating decrement scalar by multiple of 64-bit predicate constraint element count	UQDECD (scalar)
	Unsigned saturating decrement vector by multiple of 64-bit predicate constraint element count	UQDECD (vector)
UQDECP	Unsigned saturating decrement scalar by the number of predicate elements that are TRUE	UQDECP (scalar)

Mnemonic	Instruction	See
	Unsigned saturating decrement vector by the number of predicate elements that are TRUE	UQDECP (vector)
UQINCB	Unsigned saturating increment scalar by multiple of 8-bit predicate constraint element count	UQINCB
UQINCH	Unsigned saturating increment scalar by multiple of 16-bit predicate constraint element count	UQINCH (scalar)
	Unsigned saturating increment vector by multiple of 16-bit predicate constraint element count	UQINCH (vector)
UQINCW	Unsigned saturating increment scalar by multiple of 32-bit predicate constraint element count	UQINCW (scalar)
	Unsigned saturating increment vector by multiple of 32-bit predicate constraint element count	UQINCW (vector)
UQINCD	Unsigned saturating increment scalar by multiple of 64-bit predicate constraint element count	UQINCD (scalar)
	Unsigned saturating increment vector by multiple of 64-bit predicate constraint element count	UQINCD (vector)
UQINCP	Unsigned saturating increment scalar by the number of predicate elements that are TRUE	UQINCP (scalar)
	Unsigned saturating increment vector by the number of predicate elements that are TRUE	UQINCP (vector)

6.2.6.6 Loop control

I_{CSCJP}

These instructions control counted vector loops and vector loops with data-dependent termination conditions.

These instructions create a loop partition predicate with *Active elements* set to TRUE up to the point where the loop should terminate, and FALSE thereafter. Two loop concepts are supported, simple loops and data-dependent loops.

6.2.6.6.1 Simple loops

I_{FXMJD}

An up-counting WHILE instruction that increments the value of the first scalar operand and compares the value with a second, fixed scalar operand. The instruction generates a destination predicate with all of the following characteristics:

- The predicate elements starting from the lowest numbered element are true while the comparison is true.
- The predicate elements thereafter, up to the highest numbered element, are false when the comparison becomes false.

All 32 bits or 64 bits of the scalar operands are significant for the purposes of comparison. The full 32-bit or 64-bit value of the first operand is incremented by 1 for each destination predicate element, irrespective of the element size. The first general-purpose register operand is not updated.

If all of the following occur, a comparison can never fail, resulting in an all-true predicate:

- The comparison includes an equality test.
- The second scalar operand is equal to the maximum signed integer value of the selected size.

The N, Z, C, and V condition flags are unconditionally set to control a subsequent conditional branch.

Mnemonic	Instruction	See
WHILELE	While incrementing signed scalar less than or equal to scalar	WHILELE
WHILELO	While incrementing unsigned scalar lower than scalar	WHILELO
WHILELS	While incrementing unsigned scalar lower than or the same as scalar	WHILELS
WHILELT	While incrementing signed scalar less than scalar	WHILELT

6.2.6.6.2 Data-dependent loops

I_{NCXLD}

For data-dependent termination conditions, it is necessary to convert the result of a vector comparison into a loop partition predicate. The new partition truncates the current vector partition immediately before or after the first active TRUE comparison. The N, Z, C, and V condition flags are optionally set to control a subsequent conditional branch.

The BRKA instructions set active destination predicate elements to TRUE up to and including the first active TRUE element in their source predicate register, setting subsequent elements to FALSE.

The BRKB instructions set active destination predicate elements to TRUE up to but excluding the first active TRUE element in their source predicate register, setting subsequent elements to FALSE.

The BRKPA and BRKPB instructions propagate the result of a previous BRKB or BRKPB instruction, by setting their destination predicate register to all FALSE if the *Last active element* of their first source predicate register is not TRUE, but otherwise generate the destination predicate from their second source predicate as described for the BRKA and BRKB instructions.

The BRKN instructions propagate the result of a previous BRKB or BRKPB instruction by setting the destination predicate register to all FALSE if the *Last active element* of their first source predicate register is not TRUE, but otherwise leave the destination predicate unchanged. The destination and second source predicate must have been created by another instruction, such as RDFFR or WHILE.

Mnemonic	Instruction	See
BRKA	Break after the first true condition	BRKA, BRKAS
BRKAS	Break after the first true condition, setting the condition flags	BRKA, BRKAS
BRKB	Break before the first true condition	BRKB, BRKBS
BRKBS	Break before the first true condition, setting the condition flags	BRKB, BRKBS
BRKN	Propagate break to next partition	BRKN, BRKNS
BRKNS	Propagate break to next partition, setting the condition flags	BRKN, BRKNS
BRKPA	Break after the first true condition, propagating from previous partition	BRKPA, BRKPAS
BRKPAS	Break after the first true condition, propagating from previous partition, setting the condition flags	BRKPA, BRKPAS
BRKPB	Break before the first true condition, propagating from the previous partition	BRKPB, BRKPBS
BRKPBS	Break before the first true condition, propagating from the previous partition, setting the condition flags	BRKPB, BRKPBS

6.2.6.7 Serialized operations

‡_{ZVHQH}

These instructions permit *Active elements* within a vector to be processed sequentially without unpacking the vector. The condition flags are unconditionally set to control a subsequent conditional branch.

Mnemonic	Instruction	See
PFIRST	Set the <i>First active element</i> to TRUE	PFIRST
PNEXT	Find next <i>Active element</i>	PNEXT
CTERMEQ	Compare and terminate loop when equal	CTERMEQ, CTERMNE
CTERMNE	Compare and terminate loop when not equal	CTERMEQ, CTERMNE

6.2.7 Move operations

6.2.7.1 Element permute and shuffle

‡_{DNKML}

These instructions move data between different vector elements, or between vector elements and scalar registers.

These instructions perform the following operations:

- Conditionally extract the *Last active element* of a vector or the following element.
 - The supported instructions are: CLASTA, CLASTB.
- Unconditionally extract the *Last active element* of a vector or the following element.
 - The supported instructions are: LASTA, LASTB.
- Variable permute instructions where the permutation is determined by the values in a predicate register or a table of element index values.
 - The supported instructions are: COMPACT, SPLICE, TBL.
- Fixed permute instructions where the form of the permutation is encoded in the instruction.
 - The supported instructions are: DUP, EXT, INSR, REV, REVB, REVH, REVW, SUNPKHI, SUNPKLO, TRN1, TRN2, UUNPKHI, UUNPKLO, UZP1, UZP2, ZIP1, ZIP2.

Mnemonic	Instruction	See
CLASTA	Conditionally extract element after the <i>Last active element</i> to general-purpose register	CLASTA (scalar)
	Conditionally extract element after the <i>Last active element</i> to SIMD&FP scalar	CLASTA (SIMD&FP scalar)
	Conditionally extract element after the <i>Last active element</i> to vector	CLASTA (vectors)
CLASTB	Conditionally extract <i>Last active element</i> to general-purpose register	CLASTB (scalar)
	Conditionally extract <i>Last active element</i> to SIMD&FP scalar	CLASTB (SIMD&FP scalar)
	Conditionally extract <i>Last active element</i> to vector	CLASTB (vectors)
LASTA	Extract element after the <i>Last active element</i> to general-purpose register	LASTA (scalar)
	Extract element after the <i>Last active element</i> to SIMD&FP scalar	LASTA (SIMD&FP scalar)

Mnemonic	Instruction	See
LASTB	Extract <i>Last active element</i> to general-purpose register	LASTB (scalar)
	Extract <i>Last active element</i> to SIMD&FP scalar	LASTB (SIMD&FP scalar)
COMPACT	Shuffle <i>Active elements</i> of vector to the right and fill with zeros	COMPACT
SPLICE	Splice two vectors under predicate control	SPLICE
TBL	Programmable table lookup using vector of element indexes	TBL
DUP	Broadcast indexed vector element	DUP
EXT	Extract vector from pair of vectors	EXT
INSR	Insert general-purpose register into shifted vector	INSR (scalar)
	Insert SIMD&FP scalar register into shifted vector	INSR (SIMD&FP scalar)
MOV	Move indexed element or SIMD&FP scalar to vector (unpredicated)	MOV (SIMD&FP scalar, unpredicated)
	Move SIMD&FP scalar register to vector elements (predicated)	MOV (SIMD&FP scalar, predicated)
REV	Reverse all elements in vector	REV (vector)
REVB	Reverse 8-bit bytes in elements	REVB, REVH, REVW
REVH	Reverse 16-bit halfwords in elements	REVB, REVH, REVW
REVW	Reverse 32-bit words in elements	REVB, REVH, REVW
TRN1	Interleave even elements from two vectors	TRN1, TRN2 (vectors)
TRN2	Interleave odd elements from two vectors	TRN1, TRN2 (vectors)
UZP1	Concatenate even elements from two vectors	UZP1, UZP2 (vectors)
UZP2	Concatenate odd elements from two vectors	UZP1, UZP2 (vectors)
ZIP1	Interleave elements from low halves of two vectors	ZIP1, ZIP2 (vectors)
ZIP2	Interleave elements from high halves of two vectors	ZIP1, ZIP2 (vectors)

6.2.7.2 Unpacking instructions

I_{FKTHW}

These instructions unpack half of the elements from the source vector register or predicate register, widen the unpacked elements to twice the width, and place the result in the destination register.

Mnemonic	Instruction	See
SUNPKHI	Unpack and sign-extend elements from high half of vector	SUNPKHI, SUNPKLO
SUNPKLO	Unpack and sign-extend elements from low half of vector	SUNPKHI, SUNPKLO
UUNPKHI	Unpack and zero-extend elements from high half of vector	UUNPKHI, UUNPKLO
UUNPKLO	Unpack and zero-extend elements from low half of vector	UUNPKHI, UUNPKLO

Mnemonic	Instruction	See
PUNPKHI	Unpack and widen elements from high half of predicate	PUNPKHI, PUNPKLO
PUNPKLO	Unpack and widen elements from low half of predicate	PUNPKHI, PUNPKLO

6.2.7.3 Predicate permute

I_{QWFQC} These instructions are used to move and permute predicate elements. These instructions generally mirror the fixed vector permutes to allow predicates to follow their data. The permutes move all of the bits in a predicate element, not just the canonical bits.

Mnemonic	Instruction	See
REV	Reverse all elements in predicate	REV
TRN1	Interleave even elements from two predicates	TRN1, TRN2 (predicates)
TRN2	Interleave odd elements from two predicates	TRN1, TRN2 (predicates)
UZP1	Select even elements from two predicates	UZP1, UZP2 (predicates)
UZP2	Select odd elements from two predicates	UZP1, UZP2 (predicates)
ZIP1	Interleave elements from low halves of two predicates	ZIP1, ZIP2 (predicates)
ZIP2	Interleave elements from high halves of two predicates	ZIP1, ZIP2 (predicates)

6.2.7.4 Index vector generation

I_{XYTKN} The INDEX instruction initializes a vector horizontally by setting its first element to an integer value, and then repeatedly incrementing it by a second integer value to generate the subsequent elements. Each integer value can be specified as a signed immediate or a general-purpose register.

Mnemonic	Instruction	See
INDEX	Create index vector starting from and incremented by immediates	INDEX (immediates)
	Create index vector starting from immediate and incremented by general-purpose register	INDEX (immediate, scalar)
	Create index vector starting from general-purpose register and incremented by immediate	INDEX (scalar, immediate)
	Create index vector starting from and incremented by general-purpose registers	INDEX (scalars)

6.2.7.5 Move prefix

I_{TFSHM} The MOVPRFX (predicated) instruction is a predicated vector move that can be combined with a predicated destructive instruction that immediately follows it, in program order, to create a single constructive operation, or to convert an instruction with merging predication to use zeroing predication.

The MOVPRFX (unpredicated) instruction is an unpredicated vector move that can be combined with a predicated or unpredicated destructive instruction that immediately follows it, in program order, to create a single constructive operation.

The `Operational information` section of an SVE instruction description indicates whether or not an instruction can be predictably prefixed by a MOVPRFX instruction. If the `Operational information` of an SVE instruction description does not mention MOVPRFX or if the section does not exist, then the instruction cannot be predictably prefixed by a MOVPRFX instruction.

The prefixed instruction that immediately follows a MOVPRFX instruction in program order must be an SVE instruction that can be predictably prefixed by a MOVPRFX instruction, or an A64 HLT instruction, or an A64 BRK instruction. For an SVE instruction that can be predictably prefixed by a MOVPRFX instruction, all of the following apply:

- The destination register field implicitly specifies one of the source operands, which means that it is a destructive binary or ternary vector operation or unary operation with merging predication, excluding MOVPRFX.
- The destination register is the same as the MOVPRFX destination register.
- The prefixed instruction does not use the MOVPRFX destination register in any of its other source register fields, even if it has a different name but refers to the same architectural register state. For example, Z1, V1, and D1 all refer to the same architectural register.
- If the MOVPRFX instruction is predicated, then the prefixed instruction is predicated using the same Governing predicate register, and the maximum encoded element size is the same as the MOVPRFX element size, excluding the fixed-size 64-bit elements of the wide elements form of bitwise shift and integer compare operations.
- If the MOVPRFX instruction is unpredicated, then the prefixed instruction can use any Governing predicate register and element size, or it can be unpredicated. A predicated MOVPRFX cannot be used with an unpredicated instruction.

If the instruction that follows a MOVPRFX instruction is not an SVE instruction that can be predictably prefixed by a MOVPRFX instruction, the two instructions behave in one of the following `CONSTRAINED UNPREDICTABLE` ways:

- Either or both instructions can execute with their individually described effects.
- Either instruction can generate an Undefined Instruction exception.
- Either or both instructions can execute as a NOP.
- The second instruction can execute with an UNKNOWN value for any of its source registers.
- Any register that is written by either or both instructions can be set to an UNKNOWN value.
- A control flow instruction that writes the PC can set the PC to an UNKNOWN value.

Mnemonic	Instruction	See
MOVPRFX	Move prefix (predicated)	MOVPRFX
	Move prefix (unpredicated)	MOVPRFX

Unless the combination of a constructive operation with merging predication is specifically required, it is strongly recommended that, for performance reasons, software should prefer to use the zeroing form of predicated MOVPRFX or the unpredicated MOVPRFX instruction.

6.2.8 Reduction operations

6.2.8.1 Horizontal reductions

‡_{PW}NDL

These instructions perform arithmetic horizontally across *Active elements* of a single source vector and deliver a scalar result.

The floating-point horizontal accumulating sum instruction, FADDA, operates strictly in order of increasing element number across a vector, using the scalar destination register as a source for the initial value of the accumulator. This preserves the original program evaluation order where non-associativity is required.

The other floating-point reductions calculate their result using a recursive pair-wise algorithm that does not preserve the original program order, but permits increased parallelism for code that does not require strict order of evaluation.

Integer reductions are fully associative, and the order of evaluation is not specified by the architecture.

Mnemonic	Instruction	See
ANDV	Bitwise AND reduction, treating <i>Inactive elements</i> as all ones	ANDV
EORV	Bitwise XOR reduction, treating <i>Inactive elements</i> as zero	EORV
FADDA	Floating-point add strictly-ordered reduction, accumulating in scalar, ignoring <i>Inactive elements</i>	FADDA
FADDV	Floating-point add recursive reduction, treating <i>Inactive elements</i> as +0.0	FADDV
FMAXNMV	Floating-point maximum number recursive reduction, treating <i>Inactive elements</i> as the default NaN	FMAXNMV
FMAXV	Floating-point maximum recursive reduction, treating <i>Inactive elements</i> as negative infinity	FMAXV
FMINNMV	Floating-point minimum number recursive reduction, treating <i>Inactive elements</i> as the default NaN	FMINNMV
FMINV	Floating-point minimum recursive reduction, treating <i>Inactive elements</i> as positive infinity	FMINV
ORV	Bitwise OR reduction, treating <i>Inactive elements</i> as zero	ORV
SADDV	Signed add reduction, treating <i>Inactive elements</i> as zero	SADDV
SMAXV	Signed maximum reduction, treating <i>Inactive elements</i> as the minimum signed integer	SMAXV
SMINV	Signed minimum reduction, treating <i>Inactive elements</i> the maximum signed integer	SMINV
UADDV	Unsigned add reduction, treating <i>Inactive elements</i> as zero	UADDV
UMAXV	Unsigned maximum reduction, treating <i>Inactive elements</i> as zero	UMAXV
UMINV	Unsigned minimum reduction, treating <i>Inactive elements</i> as the maximum unsigned integer	UMINV

Chapter 7

SVE System and control registers

7.1 SVE modifications to the AArch64 System registers

RSNSKG The *SVE* modifications to the System Registers are all included in the following table:

Register	Change	Description
ID_AA64PFR0_EL1	Defines bits[35:32] as the <i>SVE</i> field.	Indicates whether <i>SVE</i> is implemented.
CPACR_EL1	Defines bits[17:16] as the ZEN field.	Enables access to <i>SVE</i> functionality from EL1 and EL0.
CPTR_EL2	If HCR_EL2.E2H is 0, defines bit[8] as TZ. If HCR_EL2.E2H is 1, defines bits[17:16] as the ZEN field.	The TZ field traps access to <i>SVE</i> functionality from EL2 and Non-secure EL1&0 to EL2. The ZEN field enables access to <i>SVE</i> functionality from EL2 and Non-secure EL1&0.
CPTR_EL3	Defines bit[8] as EZ.	Enables access to <i>SVE</i> functionality from EL0, EL1, EL2, and EL3.

Register	Change	Description
TCR_EL1	Defines bit[54] as NFD1 and bit[53] as NFD0 .	Disable stage 1 translation table walks caused by certain elements of the <i>SVE</i> First-fault and Non-fault vector load instructions from EL0 for translations using TTBR1_EL1 or TTBR0_EL1, respectively.
TCR_EL2	When HCR_EL2 is 1, defines bit[54] as NFD1 and bit[53] as NFD0	Disable stage 1 translation table walks caused by certain elements of the <i>SVE</i> First-fault and Non-fault vector load instructions from EL0 for translations using TTBR1_EL2 or TTBR0_EL2, respectively.
EDPFR	Defines bits[35:32] as the <i>SVE</i> field.	Indicates whether <i>SVE</i> is implemented for debug.
ESR_ELx	New exception class 0b011001 added to the EC field description.	Identifies accesses to <i>SVE</i> functionality when disabled or trapped by CPACR_EL1, CPTR_EL2, or CPTR_EL3.

See also:

- ESR_EL1
- ESR_EL2
- ESR_EL3

Draft

7.2 SVE-specific AArch64 System registers or fields

R_{ZCGZH} The SVE-specific AArch64 System registers and fields are all included in the following table:

Register	Description
ID_AA64ZFR0_EL1	SVE feature ID register 0.
ZCR_EL1	SVE control register to constrain the vector length at EL1 and EL0.
ZCR_EL2	SVE control register to constrain the vector length at EL2 and Non-secure EL1 and EL0 .
ZCR_EL3	SVE control register to constrain the vector length at EL3, EL2, EL1, and EL0.

See also:

- [4.3 Configurable vector length](#)

Draft

Chapter 8

SVE Debug

Draft

8.1 Self-hosted debug

I_{XLMZB} SVE extends the AArch64 self-hosted debug exception model.

8.1.1 SVE Watchpoint exceptions

- R_{NMHSK} For SVE predicated vector store instructions, when the instruction performs an access due to an *Active element*, a *Watchpoint debug event* is generated.
- R_{XMHCB} For SVE predicated vector load instructions that are not a First-fault or Non-fault load, when the instruction performs an access due to an *Active element*, a *Watchpoint debug event* is generated.
- R_{FZPKM} For SVE First-fault vector load instructions, when the instruction performs an access due to a *First active element*, a *Watchpoint debug event* is generated.
- R_{STCYM} For SVE unpredicated load or store instructions, when the instruction performs a byte access matching a configured watchpoint, a *Watchpoint debug event* is generated.
- R_{ZRWZT} For SVE predicated vector load or store instructions which are not First-fault vector loads or Non-fault vector loads, when the instruction performs a non-speculative single-copy atomic access matching a configured watchpoint due to an *Active element*, a *Watchpoint debug event* is generated.
- R_{GLRCD} For SVE predicated vector load or store instructions, when the instruction performs an access due to an *Inactive element*, a *Watchpoint debug event* is not generated.
- R_{DYGMS} For SVE Non-fault vector load instructions, when the instruction performs an access, a *Watchpoint debug event* is not generated.

R _{LKKQG}	For SVE Non-fault vector load instructions, when the instruction performs a non-speculative single-copy atomic access matching a configured watchpoint due to an <i>Active element</i> , the access is suppressed and reported in the FFR.
R _{TLSCX}	When all of the following are true, for SVE First-fault vector load instructions, a <i>Watchpoint debug event</i> is generated: <ul style="list-style-type: none"> • The instruction performs a non-speculative single-copy atomic access. • The access matches a configured watchpoint. • The access is for the <i>First active element</i>.
R _{XBLLW}	For self-hosted debug, the <i>Active element</i> of a First-fault vector load behaves like a regular vector load. For all subsequent elements of a First-fault vector load, the elements behave like a Non-fault vector load.
R _{CKZFP}	Watchpoints do not prevent access to memory.
R _{ZHXGG}	For SVE Non-fault and First-fault vector load instructions that do not trigger a watchpoint, an access that matches a configured watchpoint can return the data and set the corresponding FFR element to FALSE.
	See also: <ul style="list-style-type: none"> • 4.2 Synchronous memory faults

8.1.2 MOVPRFX instruction behavior in self-hosted debug

I _{JQLHW}	MOVPRFX (predicated) and MOVPRFX (unpredicated) instructions have predictable behavior when they are used with breakpoints and single-step execution.
I _{JZVZB}	A MOVPRFX instruction can be used to prefix a BRK or HLT instruction.
R _{BHPGY}	If a hardware breakpoint is programmed with the address of the prefixed instruction, a <i>Breakpoint exception</i> is generated.
R _{WLYTJ}	A hardware breakpoint is only predictable if it is programmed with the address of the initial MOVPRFX instruction, and not the address of the prefixed instruction.
R _{BMYNV}	If an instruction to be stepped uses MOVPRFX, it is CONSTRAINED UNPREDICTABLE whether a single step can either step over the pair of instructions, or step over only the MOVPRFX instruction.
	See also: <ul style="list-style-type: none"> • MOVPRFX (predicated) • MOVPRFX (unpredicated) • 4.1 Exception model

8.2 External debug

R_{NFJLQ} If the *PE* is in *Debug state*, the *SVE* architectural state can be accessed.

See also:

- *External Debug* in the *ARM® Architecture Reference Manual, ARMv8-A, for ARMv8-A architecture profile*

8.2.1 Instructions in Debug state

R_{TGSHS} If the *PE* is in *Debug state*, all of the following apply:

- The following *SVE* instructions have the same behavior as in Non-debug state:
 - RDVL.
 - CPY (immediate, zeroing) with byte element size and a shift amount of 0.
 - PTRUE, PTRUES with ALL constraint and byte element size.
 - RDIFFR (unpredicated).
 - WRFFR.
 - EXT.
 - INSR (scalar).
 - DUP (scalar).
- CMPNE (immediate) with byte element size has the same behavior as in Non-debug state, but also sets DLR_EL0 and DSPSR_EL0 to UNKNOWN values.
- For *SVE* instructions not listed in the first two bullets, their behavior is CONstrained UNPREDICTABLE. The behaviors which can occur are:
 - The instruction generates an Undefined Instruction exception.
 - The instruction executes as a NOP.
 - If the instruction modifies PSTATE, it sets DLR_EL0 and DSPSR_EL0 to UNKNOWN values.
 - If the instruction reads PSTATE condition flags, it uses an UNKNOWN value for the condition flag.
 - The instruction has the same behavior as in Non-debug state.

Chapter 9

SVE Performance Monitors Extension

Draft

9.1 SVE and PMU interaction

I _{YSGRK}	The Performance Monitors use events 0x8000-0x80FF for SVE-specific PMU events.
R _{B CYCD}	At least one of the following SVE-specific PMU events is implemented: <ul style="list-style-type: none">• SVE_INST_RETIRED• SVE_INST_SPEC
I _{RLMNH}	Arm strongly recommends that SVE_INST_RETIRED is implemented.
I _{QPFJR}	Arm recommends that the PMU events listed in <i>PMU events and event numbers</i> in the <i>Arm® Architecture Reference Manual, Armv8-A</i> are implemented.
R _{TNZGL}	For SVE recommended PMU events, an implementation can modify the definition of an event to better correspond to the implementation.
R _{QQZL}	For SVE recommended PMU events, an implementation is permitted to not use any of the events.
I _{XZMTH}	Arm recommends that the SVE-modified definition of the following PMU events is implemented: <ul style="list-style-type: none">• LD_RETIRED.• ST_RETIRED.• INST_RETIRED.• UNALIGNED_LDST_RETIRED.• MEM_ACCESS.• INST_SPEC.• MEM_ACCESS_RD.• MEM_ACCESS_WR.

- UNALIGNED_LD_SPEC.
- UNALIGNED_ST_SPEC.
- UNALIGNED_LDST_SPEC.
- LD_SPEC.
- ST_SPEC.
- LDST_SPEC.

R_{JZKXC} When any of the data movement instructions, floating-point or integer instructions, or floating-point conversions instructions are executed, all of the following are true:

- It is IMPLEMENTATION DEFINED whether operations due to any of these instructions are counted as integer operations, floating-point operations or neither of these operations.
- The instructions are not permitted to be counted as both integer operations and floating-point operations.

I_{YFWRH} Unless otherwise stated, a reference to *Advanced SIMD* or *SVE* instructions refers to all the instructions listed in [Chapter 11 Instruction categories](#) under the corresponding subheadings. This includes data-processing, predicate handling, load, store, and prefetch instructions.

I_{HVKDN} *Advanced SIMD* scalar instructions refer to the following groups of *Advanced SIMD* instructions:

- Instructions that Performance Monitors count for the ARMv8 DP_SPEC event.
- Instructions that only read element[0] of their source vectors and can write a non-zero result only to element[0] of their destination vector.

R_{PCLHP} For the Performance Monitors that count cryptographic *Advanced SIMD* instructions as CRYPTO_SPEC events, it is IMPLEMENTATION DEFINED whether or not the Performance Monitors count individual cryptographic instructions as *SIMD* instructions or *Advanced SIMD* scalar instructions.

R_{HLEFD} When a MOVPRFX instruction is executed, except for events 0x807C-0x807E, it is IMPLEMENTATION DEFINED for each execution of the instruction whether or not the Performance Monitor counts the event.

R_{WWTRW} When a MOVPRFX instruction is executed because of a MOVPRFX instruction, except for events 0x807C-0x807E, it is IMPLEMENTATION DEFINED for each execution of the operation whether or not the Performance Monitor counts the event. This can vary dynamically for each execution of the same instruction.

R_{GRTZF} When an event is architecturally executed, it is IMPLEMENTATION DEFINED whether Performance Monitors count *Operations* and *Speculative executions*.

R_{WHKYT} For microarchitectural operations executed because of an architectural instruction, the following events count speculative execution on both a correct and false execution path:

- UOP_SPEC.
- ASE_UOP_SPEC.
- SVE_UOP_SPEC.
- ASE_SVE_UOP_SPEC.
- SIMD_UOP_SPEC.

See also:

- [9.2 New performance monitor events](#)
- [9.2.1 SVE-specific PMU event descriptions](#)
- [11.1 Data movement instructions](#)
- [11.4 Floating-point conversions](#)
- [11.5 Floating-point or integer instructions](#)
- [11.6 Non-SIMD SVE instructions](#)
- [11.8.1 Cryptographic \(Advanced SIMD\)](#)

9.2 New performance monitor events

R_{JRKWQ} This following are SVE-specific PMU events:

Event number	Event type	Event mnemonic
0x8002	Architectural	SVE_INST_RETIRED
0x8006	Microarchitectural	SVE_INST_SPEC

9.2.1 SVE-specific PMU event descriptions

R_{JGMHP} The 0x8002, SVE_INST_RETIRED, SVE operations speculatively executed event counts architecturally executed SVE instructions. It is IMPLEMENTATION DEFINED whether this event counts the Non-SIMD SVE instructions.

R_{ZLBNJ} the 0x8006, SVE_INST_SPEC, SVE operations speculatively executed event counts speculatively executed operations due to SVE instructions. It is IMPLEMENTATION DEFINED whether it counts operations due to Non-SIMD SVE instructions.

See also:

- [11.6 Non-SIMD SVE instructions](#)

Draft

9.3 PMU events

R_{HPFGZ} The following PMU events are extended to count operations performed as a result of the SVE instructions:

Event number	Event mnemonic	SVE clarification
0x0006	LD_RETIRED	This event counts architecturally executed SVE load instructions.
0x0007	ST_RETIRED	This event counts architecturally executed SVE store instructions.
0x0008	INST_RETIRED	This event counts architecturally executed SVE instructions. It is IMPLEMENTATION DEFINED whether MOVPRFX is counted by this event.
0x000F	UNALIGNED_LDST_RETIRED	This event counts architecturally executed SVE load and store instructions that access at least one unaligned element address that would generate an alignment fault when alignment fault checking is enabled.
0x0013	MEM_ACCESS	This event counts memory reads and writes as a result of SVE load and store instructions. The number of accesses generated by each SVE instruction is IMPLEMENTATION DEFINED.
0x001B	INST_SPEC	This event counts speculatively executed SVE operations. It is IMPLEMENTATION DEFINED whether MOVPRFX is counted by this event.
0x0066	MEM_ACCESS_RD	This event counts memory reads as a result of SVE load and store instructions. The number of accesses generated by each SVE instruction is IMPLEMENTATION DEFINED.
0x0067	MEM_ACCESS_WR	This event counts memory writes as a result of SVE load and store instructions. The number of accesses generated by each SVE instruction is IMPLEMENTATION DEFINED.
0x0068	UNALIGNED_LD_SPEC	This event counts speculatively executed SVE load operations that access at least one unaligned element address.
0x0069	UNALIGNED_ST_SPEC	This event counts speculatively executed SVE store operations that access at least one unaligned element address.
0x006A	UNALIGNED_LDST_SPEC	This event counts speculatively executed SVE load and store operations that access at least one unaligned element address.
0x0070	LD_SPEC	This event counts speculatively executed SVE load operations.
0x0071	ST_SPEC	This event counts speculatively executed SVE store operations.
0x0072	LDST_SPEC	This event counts speculatively executed SVE load and store operations.

See also:

- *Common Event Numbers in the ARM[®] Architecture Reference Manual, ARMv8-A, for ARMv8-A architecture profile*

Chapter 10

Recommended SVE-specific PMU events

Draft

I_{BQOST}

The section titled *PMU events and event numbers* in the *Arm® Architecture Reference Manual, Armv8-A, for Armv8-A architecture profile* describes the recommended architectural and microarchitectural PMU events for SVE implementations.

10.1 Interesting combinations of SVE events

10.1.1 Scalar-equivalent operations

I_{RDHPB}

The number of speculatively executed operations performed on individual scalar values, assuming that all *SVE* vector elements are active, can be determined from a pair of event counters. For example, the total number of individual floating-point operations performed can be computed as follows:

$$FP_SCALE_OPS_SPEC \times VL \div 128 + FP_FIXED_OPS_SPEC$$

A summary of these event pairs is given below. Combined multiply-add and multiply-subtract instructions are counted as two operations per element.

Operation type	Scalable operations	Fixed width operations
Floating-point operations (any precision)	FP_SCALE_OPS_SPEC	FP_FIXED_OPS_SPEC
Half-precision floating-point operations	FP_HP_SCALE_OPS_SPEC	FP_HP_FIXED_OPS_SPEC
Single-precision floating-point operations	FP_SP_SCALE_OPS_SPEC	FP_SP_FIXED_OPS_SPEC
Double-precision floating-point operation	FP_DP_SCALE_OPS_SPEC	FP_DP_FIXED_OPS_SPEC
Integer operations (any size)	INT_SCALE_OPS_SPEC	INT_FIXED_OPS_SPEC
Load/store accesses (any size)	LDST_SCALE_OPS_SPEC	LDST_FIXED_OPS_SPEC
Load accesses (any size)	LD_SCALE_OPS_SPEC	LD_FIXED_OPS_SPEC
Store accesses (any size)	ST_SCALE_OPS_SPEC	ST_FIXED_OPS_SPEC

10.1.2 Bytes loaded and stored

I_{NYJRH}

The number of bytes speculatively loaded from memory or stored to memory, assuming that all *SVE* vector elements are active, can be determined from a pair of event counters. For example, the total number of bytes loaded from memory can be computed as follows:

$$LD_SCALE_BYTES_SPEC \times VL \div 128 + LD_FIXED_BYTES_SPEC$$

A summary of the total byte count pairs is as follows:

Operation type	Scalable operations	Fixed width operations
Load/store byte count	LDST_SCALE_BYTES_SPEC	LDST_FIXED_BYTES_SPEC
Load byte count	LD_SCALE_BYTES_SPEC	LD_FIXED_BYTES_SPEC
Store byte count	ST_SCALE_BYTES_SPEC	ST_FIXED_BYTES_SPEC

10.1.3 Overall vector utilization

I_{MVTZM}

ARM does not recommend the accumulation of an active predicate population count, or predicate weight, by every predicated *SVE* instruction. However, the vector utilization can be estimated using the result used to adjust *SVE* event counters that ignore the predicate weight and one or more of the ratios of events shown in the following:

Utilization rate	Ratio
All predicates active	$\text{SVE_PRED_FULL_SPEC} \div \text{SVE_PRED_SPEC}$
Partial predicates active	$\text{SVE_PRED_PARTIAL_SPEC} \div \text{SVE_PRED_SPEC}$
No predicates active	$\text{SVE_PRED_EMPTY_SPEC} \div \text{SVE_PRED_SPEC}$

Regions of code generating a high frequency of SVE_PRED_EMPTY_SPEC events might indicate where the addition of a B.NONE conditional branch around a block of predicated code would avoid executing instructions that frequently generate no useful result.

10.1.4 Vector loop efficiency

I_{MXYQS} The effectiveness with which sequential or scalar source loops are vectorized can be estimated using ratios of the SVE_PLOOP_*_SPEC predicated loop events, as shown in the following table:

Vector loop metric	Ratio
Source level iterations per loop	$\text{SVE_PLOOP_ELTS_SPEC} \div \text{SVE_PLOOP_TERM_SPEC}$
Vectorized iterations per loop	$\text{SVE_PLOOP_TEST_SPEC} \div \text{SVE_PLOOP_TERM_SPEC}$
Parallelism per vector loop	$\text{SVE_PLOOP_ELTS_SPEC} \div \text{SVE_PLOOP_TEST_SPEC}$

Chapter 11

Instruction categories

Draft

11.1 Data movement instructions

11.1.1 Data movement (scalar)

$\mathbb{I}_{\text{GGCPB}}$ All of the following are data movement (scalar) instructions:

- FCSEL.
- FMOV (scalar, immediate).
- FMOV (general).
- FMOV (register).

11.1.2 Data movement (Advanced SIMD)

$\mathbb{I}_{\text{BMXMV}}$ All of the following are data movement (*Advanced SIMD*) instructions:

- DUP (element).
- DUP (general).
- EXT.
- FMOV (vector, immediate).
- INS (element).
- INS (general).
- SMOV.
- TBL.
- TBX.
- TRN1.

- TRN2.
- UMOV.
- UZP1.
- UZP2.
- XTN, XTN2.
- ZIP1.
- ZIP2.

11.1.3 Data movement (SVE)

I_ZFWZB

All of the following are data movement (*SVE*) instructions:

- CLASTA (scalar).
- CLASTA (SIMD&FP scalar).
- CLASTA (vectors).
- CLASTB (scalar).
- CLASTB (SIMD&FP scalar).
- CLASTB (vectors).
- COMPACT.
- CPY (scalar).
- CPY (immediate).
- DUP (scalar).
- DUP (immediate).
- EXT.
- FCPY.
- FDUP.
- INDEX (immediate, scalar).
- INDEX (immediates).
- INDEX (scalar, immediate).
- INDEX (scalars).
- INSR (scalar).
- INSR (SIMD&FP scalar).
- LASTA (scalar).
- LASTA (SIMD&FP scalar).
- LASTB (scalar).
- LASTB (SIMD&FP scalar).
- MOVPRFX (predicated).
- MOVPRFX (unpredicated).
- REV (vector).
- SEL (vectors).
- SPLICE.
- SUNPKHI, SUNPKLO.
- TBL.
- TRN1, TRN2 (vectors)
- UUNPKHI, UUNPKLO.
- UZP1, UZP2 (vectors).
- ZIP1, ZIP2 (vectors).

Draft

11.2 Integer instructions

11.2.1 Integer (scalar)

11.2.1.1 Integer uniform arithmetic (scalar)

I_{HMLWT}

All of the following are integer uniform arithmetic (scalar) instructions:

- ADC.
- ADCS.
- ADD (extended register).
- ADD (immediate).
- ADD (shifted register).
- ADDS (extended register).
- ADDS (immediate).
- ADDS (shifted register).
- CCMN (immediate).
- CCMN (register).
- CCMP (immediate).
- CCMP (register).
- CSINC.
- CSINV.
- CSNEG.
- MADD.
- MSUB.
- SBC.
- SBCS.
- SDIV.
- UDIV.
- SMULH.
- UMULH.
- SUB (extended register).
- SUB (immediate).
- SUB (shifted register).
- SUBS (extended register).
- SUBS (immediate).
- SUBS (shifted register).
- ADR.
- ADRP.

11.2.1.2 Integer widening arithmetic

I_{ZNGLP}

All of the following are integer widening arithmetic instructions:

- SMADDL.
- SMSUBL.
- UMADDL.
- UMSUBL.

11.2.1.3 Integer bitwise operations (scalar)

I_{GTHXC}

All of the following are integer bitwise operations (scalar) instructions:

- AND (immediate)

- AND (shifted register)
- ANDS (immediate)
- ANDS (shifted register)
- BIC (shifted register).
- BICS (shifted register).
- EOR (immediate).
- EOR (shifted register).
- EON (shifted register).
- ORR (immediate).
- ORR (shifted register).
- ORN (shifted register).
- ASRV.
- LSLV.
- LSRV.
- RORV.
- BFM.
- SBFM.
- UBFM.
- CLS.
- CLZ.
- EXTR.
- RBIT.
- REV.
- REV16.
- REV32.

11.2.2 Integer (Advanced SIMD)

11.2.2.1 Integer uniform arithmetic (Advanced SIMD)

I_QBHGF

All of the following are integer uniform arithmetic (*Advanced SIMD*) instructions:

- ABS.
- NEG (vector).
- ADD (vector).
- SUB (vector).
- MLA (by element).
- MLA (vector).
- MLS (by element).
- MLS (vector).
- MUL (by element).
- MUL (vector).
- PMUL.
- SABA.
- UABA.
- SABD.
- UABD.
- SDOT (by element).
- SDOT (vector).
- UDOT (by element).
- UDOT (vector).
- SHADD.
- SHSUB.
- SRHADD.

- UHADD.
- UHSUB.
- URHADD.
- SMAX.
- SMIN.
- UMAX.
- UMIN.
- SQABS.
- SQNEG.
- SQADD.
- SQSUB.
- SUQADD.
- UQADD.
- USQADD.
- UQSUB.
- SQDMULH (by element).
- SQDMULH (vector).
- SQRDMULH (by element).
- SQRDMULH (vector).
- SQRDMLAH (by element).
- SQRDMLAH (vector).
- SQRDMLSH (by element).
- SQRDMLSH (vector).
- URECPE.
- URSQRTE.
- USRA.

11.2.2.2 Integer widening arithmetic (Advanced SIMD)

I_{CGVC}

All of the following are integer widening arithmetic (*Advanced SIMD*) instructions:

- SABAL, SABAL2.
- UABAL, UABAL2.
- SABDL, SABDL2.
- UABDL, UABDL2.
- SADDL, SADDL2.
- UADDL, UADDL2.
- SADDW, SADDW2.
- UADDW, UADDW2.
- SMLAL, SMLAL2 (by element).
- SMLAL, SMLAL2 (vector).
- UMLAL, UMLAL2 (by element).
- UMLAL, UMLAL2 (vector).
- SMLSL, SMLSL2 (by element).
- SMLSL, SMLSL2 (vector).
- UMLSL, UMLSL2 (by element).
- UMLSL, UMLSL2 (vector).
- SMULL, SMULL2 (by element).
- SMULL, SMULL2 (vector).
- UMULL, UMULL2 (by element).
- UMULL, UMULL2 (vector).
- PMULL, PMULL2.
- SQDMULL, SQDMULL2 (by element).
- SQDMULL, SQDMULL2 (vector).
- SQDMLAL, SQDMLAL2 (by element).

- SQDMLAL, SQDMLAL2 (vector).
- SQDMLSL, SQDMLSL2 (by element).
- SQDMLSL, SQDMLSL2 (vector).
- SHLL, SHLL2.
- SSHLL, SSHLL2.
- USHLL, USHLL2.
- SSUBL, SSUBL2.
- USUBL, USUBL2.
- SSUBW, SSUBW2.
- USUBW, USUBW2.
- UXTL, UXTL2.

11.2.2.3 Integer narrowing arithmetic (Advanced SIMD)

I_{YFSMM}

All of the following are integer narrowing arithmetic (*Advanced SIMD*) instructions:

- ADDHN, ADDHN2.
- RADDHN, RADDHN2.
- SUBHN, SUBHN2.
- RSUBHN, RSUBHN2.
- SHRN, SHRN2.
- RSHRN, RSHRN2.
- SQSHRN, SQSHRN2.
- SQSHRUN, SQSHRUN2.
- UQRSHRN, UQRSHRN2.
- UQSHRN, UQSHRN2.
- SQXTN, SQXTN2.
- SQXTUN, SQXTUN2.
- UQXTN, UQXTN2.

Draft

11.2.2.4 Integer bitwise operations (Advanced SIMD)

I_{DCGLC}

All of the following are integer bitwise operations (*Advanced SIMD*) instructions:

- AND (vector).
- BIC (vector, immediate).
- BIC (vector, register).
- EOR (vector).
- ORN (vector).
- ORR (vector, immediate).
- ORR (vector, register).
- BIF.
- BIT.
- BSL.
- CLS (vector).
- CLZ (vector).
- CNT.
- MOVI.
- MVNI.
- NOT.
- RBIT (vector).
- REV16 (vector).
- REV32 (vector).
- REV64.
- SHL.

- SRSHL.
- URSHL.
- SRSHR.
- URSHR.
- SRSRA.
- SSRA.
- URSRA.
- SLI.
- SRI.
- SQRSHL.
- SQSHL (register).
- SQSHLU.
- UQRSHL.
- UQSHL (immediate).
- UQSHL (register).
- SSSL.
- USHL.
- SSHR.
- USHR.

11.2.2.5 Integer comparisons (Advanced SIMD)

I_{TVYLX}

All of the following are integer comparisons (*Advanced SIMD*) instructions:

- CMEQ (register).
- CMEQ (zero).
- CMGE (register).
- CMGE (zero).
- CMGT (register).
- CMGT (zero).
- CMHI (register).
- CMHS (register).
- CMLE (zero).
- CMLT (zero).
- CMTST.

11.2.2.6 Integer reductions (Advanced SIMD)

I_{SNQKR}

All of the following are integer reductions (*Advanced SIMD*) instructions:

- ADDP (scalar).
- ADDP (vector).
- ADDV (vector).
- SADALP.
- UADALP.
- SADDLP.
- SADDLV.
- UADDLP.
- UADDLV.
- SMAXP.
- SMAXV.
- UMAXP.
- UMAXV.
- SMINP.
- SMINV.

- UMINP.
- UMINV.

11.2.3 Integer (SVE)

11.2.3.1 Integer uniform arithmetic (SVE)

I_{GMVWX}

All of the following are integer uniform arithmetic (SVE) instructions:

- ABS.
- NEG.
- ADD (immediate).
- ADD (vectors, predicated).
- ADD (vectors, unpredicated).
- SUB (immediate).
- SUB (vectors, predicated).
- SUB (vectors, unpredicated).
- SUBR (immediate).
- SUBR (vectors).
- ADR.
- CNOT.
- MAD.
- MSB.
- MLA (indexed).
- MLA (vectors).
- MLS (indexed).
- MLS (vectors).
- MUL (immediate).
- MUL (indexed).
- MUL (vectors, predicated).
- MUL (vectors, unpredicated).
- SABD.
- UABD.
- SDIV.
- SDIVR.
- UDIV.
- UDIVR.
- SDOT (indexed).
- SDOT (vectors).
- UDOT (indexed).
- UDOT (vectors).
- SMAX (immediate).
- SMAX (vectors).
- SMIN (immediate).
- SMIN (vectors).
- UMAX (immediate).
- UMAX (vectors).
- UMIN (immediate).
- UMIN (vectors).
- SMULH (predicated).
- SMULH (unpredicated).
- UMULH (predicated).
- UMULH (unpredicated).
- SQADD (immediate).

Draft

- SQADD (vectors, predicated).
- SQADD (vectors, unpredicated).
- SQSUB (immediate).
- SQSUB (vectors, predicated).
- SQSUB (vectors, unpredicated).
- UQADD (immediate).
- UQADD (vectors, predicated).
- UQADD (vectors, unpredicated).
- UQSUB (immediate).
- UQSUB (vectors, predicated).
- UQSUB (vectors, unpredicated).
- SXTB, SXTH, SXTW.
- UXTB, UXTH, UXTW.

11.2.3.2 Integer bitwise operations (SVE)

‡_PNLPPF

All of the following are integer bitwise operations (SVE) instructions:

- AND (vectors, predicated).
- AND (vectors, unpredicated).
- BIC (vectors, predicated).
- BIC (vectors, unpredicated).
- EON.
- EOR (vectors, predicated).
- EOR (vectors, unpredicated).
- ORN.
- ORR (vectors, predicated).
- ORR (vectors, unpredicated).
- ASR (immediate, predicated).
- ASR (immediate, unpredicated).
- ASR (vectors).
- ASR (wide elements, predicated).
- ASR (wide elements, unpredicated).
- ASRR.
- ASRD.
- CLS.
- CLZ.
- CNT.
- DUPM.
- LSL (immediate, predicated).
- LSL (immediate, unpredicated).
- LSL (vectors).
- LSL (wide elements, predicated).
- LSL (wide elements, unpredicated).
- LSLR.
- LSR (immediate, predicated).
- LSR (immediate, unpredicated).
- LSR (vectors).
- LSR (wide elements, predicated).
- LSR (wide elements, unpredicated).
- LSRR.
- NOT (vector).
- RBIT.
- REVB, REVH, REVW.

11.2.3.3 Integer comparisons (SVE)

I_{VYSCC}

All of the following are integer comparisons (SVE) instructions:

- CMP<cc> (immediate).
- CMP<cc> (vectors).
- CMP<cc> (wide elements).

11.2.3.4 Integer reductions (SVE)

I_{THTPR}

All of the following are integer reductions (SVE) instructions:

- ANDV.
- EORV.
- ORV.
- SADDV.
- UADDV.
- SMAXV.
- UMAXV.
- SMINV.
- UMINV.

11.2.3.5 Element count and increment vector (SVE)

I_{MPWZH}

All of the following are element count and increment vector (SVE) instructions:

- DECD, DECH, DECW (vector).
- INCD, INCH, INCW (vector).
- SQDECH (vector).
- SQDECW (vector).
- SQDECD (vector).
- SQINCH (vector).
- SQINCW (vector).
- SQINCD (vector).
- UQDECH (vector).
- UQDECW (vector).
- UQDECD (vector).
- UQINCH (vector).
- UQINCW (vector).
- UQINCD (vector).

11.3 Floating-point instructions

11.3.1 Floating-point (scalar)

11.3.1.1 Floating-point arithmetic (scalar)

I_{WLHZN}

All of the following are floating-point arithmetic (scalar) instructions:

- FADD (scalar).
- FSUB (scalar).
- FDIV (scalar).
- FMADD.
- FMSUB.
- FNMADD.
- FNMSUB.
- FMUL (scalar).
- FNMUL (scalar).
- FSQRT (scalar).

11.3.1.2 Floating-point miscellaneous (scalar)

I_{DMZTD}

All of the following are floating-point miscellaneous (scalar) instructions:

- FMAX (scalar).
- FMAXNM (scalar).
- FMIN (scalar).
- FMINNM (scalar).
- FRINTA (scalar).
- FRINTI (scalar).
- FRINTM (scalar).
- FRINTN (scalar).
- FRINTP (scalar).
- FRINTX (scalar).
- FRINTZ (scalar).

11.3.1.3 Floating-point comparisons (scalar)

I_{QYXGH}

All of the following are floating-point comparisons (scalar) instructions:

- FCMP.
- FCMPE.

11.3.2 Floating-point (Advanced SIMD)

11.3.2.1 Floating-point arithmetic (Advanced SIMD)

I_{PTKPC}

All of the following are floating-point arithmetic (*Advanced SIMD*) instructions:

- FABD.
- FADD (vector).
- FSUB (vector).
- FCADD.
- FCMLA.
- FCMLA (by element).

- FDIV (vector).
- FMLA (by element).
- FMLA (vector).
- FMLS (by element).
- FMLS (vector).
- FMUL (by element).
- FMUL (vector).
- FMULX.
- FMULX (by element).
- FRECPX.
- FRSQRTS.
- FSQRT (vector).

11.3.2.2 Floating-point miscellaneous (Advanced SIMD)

I_{GJCYK}

All of the following are floating-point miscellaneous (*Advanced SIMD*) instructions:

- FMAX (vector).
- FMAXNM (vector).
- FMIN (vector).
- FMINNM (vector).
- FRECPX.
- FRINTA (vector).
- FRINTI (vector).
- FRINTM (vector).
- FRINTN (vector).
- FRINTP (vector).
- FRINTX (vector).
- FRINTZ (vector).

11.3.2.3 Floating-point comparisons (Advanced SIMD)

I_{XVJNF}

All of the following are floating-point comparisons (*Advanced SIMD*) instructions:

- FACGE.
- FACGT.
- FCMEQ (register).
- FCMEQ (zero).
- FCMGE (register).
- FCMGE (zero).
- FCMGT (register).
- FCMGT (zero).
- FCMLE (zero).
- FCMLT (zero).

11.3.2.4 Floating-point reductions (Advanced SIMD)

I_{FNRS}

All of the following are floating-point reductions (*Advanced SIMD*) instructions:

- FADDP (scalar).
- FADDP (vector).
- FMAXNMP (scalar).
- FMAXNMP (vector).
- FMAXP (scalar).
- FMAXP (vector).
- FMAXNMV.

- FMAXV.
- FMINNMP (scalar).
- FMINNMP (vector).
- FMINP (scalar).
- FMINP (vector).
- FMINNMV.
- FMINV.

11.3.3 Floating-point (SVE)

11.3.3.1 Floating-point arithmetic (SVE)

I_{BGBSK}

All of the following are floating-point arithmetic (*SVE*) instructions:

- FABD.
- FADD (immediate).
- FADD (vectors, predicated).
- FADD (vectors, unpredicated).
- FSUB (immediate).
- FSUB (vectors, predicated).
- FSUB (vectors, unpredicated).
- FSUBR (immediate).
- FSUBR (vectors).
- FCADD.
- FCMLA (indexed).
- FCMLA (vectors).
- FDIV.
- FDIVR.
- FMAD.
- FNMAD.
- FNMSB.
- FMSB.
- FMLA (indexed).
- FMLA (vectors).
- FMLS (indexed).
- FMLS (vectors).
- FNMLA.
- FNMLS.
- FMUL (immediate).
- FMUL (indexed).
- FMUL (vectors, predicated).
- FMUL (vectors, unpredicated).
- FMULX.
- FRECPS.
- FRSQRTS.
- FSCALE.
- FSQRT.
- FTMAD.
- FTSMUL.

11.3.3.2 Floating-point miscellaneous (SVE)

I_{QHKWQ}

All of the following are floating-point miscellaneous (*SVE*) instructions:

- FMAX (immediate).

- FMAX (vectors).
- FMAXNM (immediate).
- FMAXNM (vectors).
- FMIN (immediate).
- FMIN (vectors).
- FMINNM (immediate).
- FMINNM (vectors).
- FRECPX.
- FRINTA.
- FRINTI.
- FRINTM.
- FRINTN.
- FRINTP.
- FRINTX.
- FRINTZ.

11.3.3.3 Floating-point comparisons (SVE)

‡_{CCKRC}

All of the following are floating-point comparisons (*SVE*) instructions:

- FACGE.
- FACGT.
- FACLE.
- FACLT.
- FCMEQ (vectors).
- FCMEQ (zero).
- FCMGE (vectors).
- FCMGE (zero).
- FCMGT (vectors).
- FCMGT (zero).
- FCMLE (vectors).
- FCMLE (zero).
- FCMGT (vectors).
- FCMGT (zero).
- FCMLT (vectors).
- FCMLT (zero).
- FCMNE (vectors).
- FCMNE (zero).
- FCMUO (vectors).

11.3.3.4 Floating-point reductions (SVE)

‡_{RBLQR}

All of the following are floating-point reductions (*SVE*) instructions:

- FADDA.
- FADDV.
- FMAXNMV.
- FMAXV.
- FMINNMV.
- FMINV.

11.4 Floating-point conversions

11.4.1 Float↔Float convert (scalar)

I_{YVFPQ} The following is a Floating-point convert (scalar) instruction:

- FCVT.

11.4.2 Float↔Float convert (Advanced SIMD)

I_{MFPKF} All of the following are Floating-point convert (Advanced SIMD) instructions:

- FCVTL, FCVTL2.
- FCVTN, FCVTN2.
- FCVTXN, FCVTXN2.

11.4.3 Float↔Float convert (SVE)

I_{WCKFQ} The following is a Floating-point convert (SVE) instruction:

- FCVT.

11.4.4 Float↔Int convert (scalar)

I_{PCMBH} All of the following are Floating-point integer convert (scalar) instructions:

- FCVTAS (scalar).
- FCVTMS (scalar).
- FCVTNS (scalar).
- FCVTPS (scalar).
- FCVTZS (scalar, fixed-point).
- FCVTZS (scalar, integer).
- FCVTAU (scalar).
- FCVTMU (scalar).
- FCVTNU (scalar).
- FCVTPU (scalar).
- FCVTZU (scalar, fixed-point).
- FCVTZU (scalar, integer).
- FJCVTZS.
- SCVTF (scalar, fixed-point).
- SCVTF (scalar, integer).
- UCVTF (scalar, fixed-point).
- UCVTF (scalar, integer).

11.4.5 Float↔Int convert (Advanced SIMD)

I_{SGHBH} All of the following are Floating-point integer convert (Advanced SIMD) instructions:

- FCVTAS (vector).
- FCVTMS (vector).
- FCVTNS (vector).
- FCVTPS (vector).

- FCVTZS (vector, fixed-point).
- FCVTZS (vector, integer).
- FCVTZS (vector, integer).
- FCVTAU (vector).
- FCVTMU (vector).
- FCVTNU (vector).
- FCVTPU (vector).
- FCVTZU (vector, fixed-point).
- FCVTZU (vector, integer).
- SCVTF (vector, fixed-point).
- SCVTF (vector, integer).
- UCVTF (vector, fixed-point).
- UCVTF (vector, integer)

11.4.6 Float↔Int convert (SVE)

‡_{KDNBP}

All of the following are Floating-point integer convert (SVE) instructions:

- FCVTZS.
- FCVTZU.
- SCVTF.
- UCVTF.

Draft

11.5 Floating-point or integer instructions

11.5.1 Floating-point or integer arithmetic (scalar)

- \mathbb{I}_{YFHSZ} All of the following are Floating-point or integer arithmetic (scalar) instructions:
- FABS (scalar).
 - FNEG (scalar).

11.5.2 Floating-point or integer arithmetic (Advanced SIMD)

- \mathbb{I}_{JZJVJ} All of the following are Floating-point or integer arithmetic (Advanced SIMD) instructions:
- FABS (vector).
 - FNEG (vector).
 - FRECPE.
 - FRSQRTE.

11.5.3 Floating-point or integer arithmetic (SVE)

- \mathbb{I}_{XFWRK} All of the following are Floating-point or integer arithmetic (SVE) instructions:
- FABS.
 - FNEG.
 - FRECPE.
 - FRSQRTE.
 - FEXPA.
 - FTSEL.

Draft

11.6 Non-SIMD SVE instructions

11.6.1 Element count and increment scalar (SVE)

I_{GVGGT}

All of the following are element count and increment scalar (*SVE*) instructions:

- ADDPL.
- ADDPL.
- RDVL.
- CNTB, CNTD, CNTH, CNTW.
- DECB, DECD, DECH, DECW (scalar).
- INCB, INCD, INCH, INCW (scalar).
- SQDECB.
- SQDECH (scalar).
- SQDECW (scalar).
- SQDECD (scalar).
- SQINCB.
- SQINCH (scalar).
- SQINCW (scalar).
- SQINCD (scalar).
- UQDECB.
- UQDECH (scalar).
- UQDECW (scalar).
- UQDECD (scalar).
- UQINCB.
- UQINCH (scalar).
- UQINCW (scalar).
- UQINCD (scalar).

Draft

11.6.2 Compare and terminate (SVE)

I_{KHHTX}

The following is a compare and terminate (*SVE*) instruction:

- CTERM EQ, CTERM NE.

11.7 Predicate instructions

11.7.1 Predicate move (SVE)

$\mathbb{I}_{\text{HFNFD}}$ All of the following are predicate move (SVE) instructions:

- PFALSE.
- PTRUE, PTRUES.
- PUNPKHI, PUNPKLO.
- RDIFFR, RDIFFRS (predicated).
- RDIFFR (unpredicated).
- SETFFR.
- WRFFR.
- REV (predicate).
- SEL (predicates).
- TRN1, TRN2 (predicates).
- UZP1, UZP2 (predicates).
- ZIP1, ZIP2 (predicates).

11.7.2 Predicate counted loop (SVE)

$\mathbb{I}_{\text{YPLYS}}$ All of the following are predicate counted loop (SVE) instructions:

- WHILELE.
- WHILELO.
- WHILELS.
- WHILELT.

11.7.3 Predicate bitwise logical operations (SVE)

$\mathbb{I}_{\text{JGDHG}}$ All of the following are predicate bitwise logical operations (SVE) instructions:

- AND, ANDS (predicates).
- BIC, BICS (predicates).
- EOR, EORS (predicates).
- NAND, NANDS.
- NOR, NORS.
- NOT (predicate).
- NOTS.
- ORN, ORNS (predicates).
- ORR, ORRS (predicates).
- PTEST.

11.7.4 Predicate scan (SVE)

$\mathbb{I}_{\text{NKBWD}}$ All of the following are predicate scan (SVE) instructions:

- BRKA, BRKAS.
- BRKB, BRKBS.
- BRKN, BRKNS.
- BRKPA, BRKPAS.
- BRKPB, BRKPBS.
- PFIRST.

11.7.5 Predicate count and increment scalar (SVE)

‡_{JRGKK}

All of the following are predicate count and increment scalar (*SVE*) instructions:

- CNTP.
- DECP (scalar).
- INCP (scalar).
- SQDECP (scalar).
- SQINCP (scalar).
- UQDECP (scalar).
- UQINCP (scalar).

11.7.6 Predicate count and increment vector (SVE)

‡_{HQZKL}

All of the following are predicate count and increment vector (*SVE*) instructions:

- DECP (vector).
- INCP (vector).
- SQDECP (vector).
- SQINCP (vector).
- UQDECP (vector).
- UQINCP (vector).

Draft

11.8 Cryptographic instructions

11.8.1 Cryptographic (Advanced SIMD)

‡_{GBYWL}

All of the following are Cryptographic (Advanced SIMD) instructions:

- AESD.
- AESE.
- AESIMC.
- AESMC.
- SHA1C.
- SHA1H.
- SHA1M.
- SHA1P.
- SHA1SU0.
- SHA1SU1.
- SHA256H.
- SHA256H2.
- SHA256SU0.
- SHA256SU1.

Draft

11.9 Load/store/prefetch instructions

11.9.1 Load/store (Advanced SIMD and floating-point scalar)

11.9.1.1 Contiguous elements load/store (Advanced SIMD)

I_{VRXTK} All of the following are contiguous elements load/store (*Advanced SIMD*) instructions:

- LD1 (multiple structures).
- ST1 (multiple structures).

11.9.1.2 Contiguous structures load/store (Advanced SIMD)

I_{LZRDO} All of the following are contiguous structures load/store (*Advanced SIMD*) instructions:

- LD2 (multiple structures).
- LD3 (multiple structures).
- LD4 (multiple structures).
- ST2 (multiple structures).
- ST3 (multiple structures).
- ST4 (multiple structures).

11.9.1.3 Single element/structure load/store (Advanced SIMD)

I_{KLWGR} All of the following are single element/structure load/store (*Advanced SIMD*) instructions:

- LD1 (single structure).
- LD2 (single structure).
- LD3 (single structure).
- LD4 (single structure).
- ST1 (single structure).
- ST2 (single structure).
- ST3 (single structure).
- ST4 (single structure).

11.9.1.4 Single element/structure replicating load (Advanced SIMD)

I_{XQCVF} All of the following are single element/structure replicating load (*Advanced SIMD*) instructions:

- LD1R.
- LD2R.
- LD3R.
- LD4R.

11.9.1.5 Register load/store (Advanced SIMD and floating-point scalar)

I_{LSJXF} All of the following are register load/store (*Advanced SIMD&FP* scalar) instructions:

- LDNP (SIMD&FP).
- LDP (SIMD&FP).
- LDR (immediate, SIMD&FP).
- LDR (literal, SIMD&FP).
- LDR (register, SIMD&FP).
- LDUR (SIMD&FP).
- STNP (SIMD&FP).

- STP (SIMD&FP).
- STR (immediate, SIMD&FP).
- STR (register, SIMD&FP).
- STUR (SIMD&FP).

11.9.2 Load/store/prefetch (SVE)

11.9.2.1 Contiguous elements load/store/prefetch (SVE)

I_{NNZKF}

All of the following are contiguous elements load/store/prefetch (SVE) instructions:

- LD1B (scalar plus immediate).
- LD1H (scalar plus immediate).
- LD1W (scalar plus immediate).
- LD1D (scalar plus immediate).
- LD1SB (scalar plus immediate).
- LD1SH (scalar plus immediate).
- LD1SW (scalar plus immediate).
- LD1B (scalar plus scalar).
- LD1H (scalar plus scalar).
- LD1W (scalar plus scalar).
- LD1D (scalar plus scalar).
- LD1SB (scalar plus scalar).
- LD1SH (scalar plus scalar).
- LD1SW (scalar plus scalar).
- LDFF1B (scalar plus scalar).
- LDFF1H (scalar plus scalar).
- LDFF1W (scalar plus scalar).
- LDFF1D (scalar plus scalar).
- LDFF1SB (scalar plus scalar).
- LDFF1SH (scalar plus scalar).
- LDFF1SW (scalar plus scalar).
- LDNF1B.
- LDNF1H.
- LDNF1W.
- LDNF1D.
- LDNF1SB.
- LDNF1SH.
- LDNT1B (scalar plus immediate).
- LDNT1H (scalar plus immediate).
- LDNT1W (scalar plus immediate).
- LDNT1D (scalar plus immediate).
- LDNT1B (scalar plus scalar).
- LDNT1H (scalar plus scalar).
- LDNT1W (scalar plus scalar).
- LDNT1D (scalar plus scalar).
- PRFB (scalar plus immediate).
- PRFH (scalar plus immediate).
- PRFW (scalar plus immediate).
- PRFH (scalar plus immediate).
- PRFB (scalar plus scalar).
- PRFH (scalar plus scalar).
- PRFW (scalar plus scalar).
- PRFD (scalar plus scalar).

- ST1B (scalar plus immediate).
- ST1H (scalar plus immediate).
- ST1W (scalar plus immediate).
- ST1D (scalar plus immediate).
- ST1B (scalar plus scalar).
- ST1H (scalar plus scalar).
- ST1W (scalar plus scalar).
- ST1D (scalar plus scalar).
- STNT1B (scalar plus immediate).
- STNT1H (scalar plus immediate).
- STNT1W (scalar plus immediate).
- STNT1D (scalar plus immediate).
- STNT1B (scalar plus scalar).
- STNT1H (scalar plus scalar).
- STNT1W (scalar plus scalar).
- STNT1D (scalar plus scalar).

11.9.2.2 Contiguous structures load/store (SVE)

↳ QPNSV

All of the following are contiguous structures load/store (SVE) instructions:

- LD2B (scalar plus immediate).
- LD2H (scalar plus immediate).
- LD2W (scalar plus immediate).
- LD2D (scalar plus immediate).
- LD2B (scalar plus scalar).
- LD2H (scalar plus scalar).
- LD2W (scalar plus scalar).
- LD2D (scalar plus scalar).
- LD3B (scalar plus immediate).
- LD3H (scalar plus immediate).
- LD3W (scalar plus immediate).
- LD3D (scalar plus immediate).
- LD3B (scalar plus scalar).
- LD3H (scalar plus scalar).
- LD3W (scalar plus scalar).
- LD3D (scalar plus scalar).
- LD4B (scalar plus immediate).
- LD4H (scalar plus immediate).
- LD4W (scalar plus immediate).
- LD4D (scalar plus immediate).
- LD4B (scalar plus scalar).
- LD4H (scalar plus scalar).
- LD4W (scalar plus scalar).
- LD4D (scalar plus scalar).
- ST2B (scalar plus immediate).
- ST2H (scalar plus immediate).
- ST2W (scalar plus immediate).
- ST2D (scalar plus immediate).
- ST2B (scalar plus scalar).
- ST2H (scalar plus scalar).
- ST2W (scalar plus scalar).
- ST2D (scalar plus scalar).
- ST3B (scalar plus immediate).
- ST3H (scalar plus immediate).

- ST3W (scalar plus immediate).
- ST3D (scalar plus immediate).
- ST3B (scalar plus scalar).
- ST3H (scalar plus scalar).
- ST3W (scalar plus scalar).
- ST3D (scalar plus scalar).
- ST4B (scalar plus immediate).
- ST4H (scalar plus immediate).
- ST4W (scalar plus immediate).
- ST4D (scalar plus immediate).
- ST4B (scalar plus scalar).
- ST4H (scalar plus scalar).
- ST4W (scalar plus scalar).
- ST4D (scalar plus scalar).

11.9.2.3 Gather/scatter load/store/prefetch (SVE)

‡CDKPPW

All of the following are gather/scatter load/store/prefetch (*SVE*) instructions:

- LD1B (vector plus immediate).
- LD1H (vector plus immediate).
- LD1W (vector plus immediate).
- LD1D (vector plus immediate).
- LD1SB (vector plus immediate).
- LD1SH (vector plus immediate).
- LD1SW (vector plus immediate).
- LD1B (scalar plus vector).
- LD1H (scalar plus vector).
- LD1W (scalar plus vector).
- LD1D (scalar plus vector).
- LD1SB (scalar plus vector).
- LD1SH (scalar plus vector).
- LD1SW (scalar plus vector).
- LDFF1B (vector plus immediate).
- LDFF1H (vector plus immediate).
- LDFF1W (vector plus immediate).
- LDFF1D (vector plus immediate).
- LDFF1SB (vector plus immediate).
- LDFF1SH (vector plus immediate).
- LDFF1SW (vector plus immediate).
- LDFF1B (scalar plus vector).
- LDFF1H (scalar plus vector).
- LDFF1W (scalar plus vector).
- LDFF1D (scalar plus vector).
- LDFF1SB (scalar plus vector).
- LDFF1SH (scalar plus vector).
- LDFF1SW (scalar plus vector).
- PRFB (vector plus immediate).
- PRFH (vector plus immediate).
- PRFW (vector plus immediate).
- PRFD (vector plus immediate).
- PRFB (scalar plus vector).
- PRFH (scalar plus vector).
- PRFW (scalar plus vector).
- PRFD (scalar plus vector).

- ST1B (vector plus immediate).
- ST1H (vector plus immediate).
- ST1W (vector plus immediate).
- ST1D (vector plus immediate).
- ST1B (scalar plus vector).
- ST1H (scalar plus vector).
- ST1W (scalar plus vector).
- ST1D (scalar plus vector).

11.9.2.4 Single element load and replicate (SVE)

I_{WRTF}

All of the following are single element load and replicate (*SVE*) instructions:

- LD1RB.
- LD1RH.
- LD1RW.
- LD1RD.
- LD1RSB.
- LD1RSH.
- LD1RSW.

11.9.2.5 Single quadword load and replicate (SVE)

I_{SRJYM}

All of the following are single quadword load and replicate (*SVE*) instructions:

- LD1RQB (scalar plus immediate).
- LD1RQH (scalar plus immediate).
- LD1RQW (scalar plus immediate).
- LD1RQD (scalar plus immediate).
- LD1RQB (scalar plus scalar).
- LD1RQH (scalar plus scalar).
- LD1RQW (scalar plus scalar).
- LD1RQD (scalar plus scalar).

11.9.2.6 Register load/store (SVE)

I_{GHZLB}

All of the following are register load/store (*SVE*) instructions:

- LDR (predicate).
- LDR (vector).
- STR (predicate).

Chapter 12

Glossary

Draft

Active element

An Active element is a vector element or predicate element that is a source register element or destination register element used by an instruction. When the corresponding element in the instruction's Governing predicate is TRUE, the element is Active. If an instruction is unpredicated, all of the vector elements or predicate elements are implicitly Active.

Advanced SIMD instructions

It is IMPLEMENTATION DEFINED if this includes the *Advanced SIMD* cryptographic instructions that would be counted by the Armv8 CRYPTO_SPEC event.

Advanced SIMD scalar instructions

A reference either to instructions that would be counted for the Armv8 DP_SPEC event, or to *Advanced SIMD* instructions which only read element[0] of their source vectors and can write a non-zero result only to element[0] of their destination vector.

Constructive instruction encoding

A constructive instruction encoding is an instruction encoding where the destination register is encoded independently of the source registers.

Destructive instruction encoding

A destructive instruction encoding is an instruction encoding where one of the source registers is also used as the destination register.

Element number

For a given element size of N bits, elements within a vector or predicate register are numbered with element[0] always representing bits[(N-1):0], element[1] always representing bits[(2N-1):N], and so on. For more information, see the layout diagram in [2.1.1 SVE Vector registers](#).

First active element

The First active element of a vector or predicate register is the lowest numbered element that is an Active element.

First-fault load

SVE provides a First-fault option for some *SVE* vector load instructions. This option causes memory access faults to be suppressed if they do not occur as a result of the First active element of the vector. Instead, the FFR is updated to indicate which of the active vector elements were not successfully loaded. For more information, see [2.1.3 First Fault Register, FFR](#).

Gather-load

Gather-load is a mechanism that allows the elements of a vector to be read from non-contiguous memory locations using a vector of addresses, where the addresses are constructed according to the addressing mode. For more information, see [11.9 Load/store/prefetch instructions](#).

Governing predicate

The predicate register that is used to determine the Active elements of a predicated instruction is known as the Governing predicate for that instruction.

Inactive element

An Inactive element is a vector element or predicate element that is an unused source register element or destination register element for the associated instruction. When the corresponding element of an instruction's Governing predicate is FALSE, the element is inactive.

Last active element

The Last active element of a vector or predicate register is the highest numbered element that is an Active element.

Memory element

An item of data in memory that is transferred to or from a vector or predicate element by an *SVE* load or store instruction. Each memory element has an access size and a type. The memory element access size is specified by each load and store instruction independently of the vector element size.

Merging predication

When a predicated instruction specifies merging predication, the Inactive elements of the destination register remain unchanged.

Non-fault load

SVE provides a Non-fault option for some *SVE* vector load instructions. This option causes all memory access faults to be suppressed. Instead, the FFR is updated to indicate which of the active vector elements were not successfully loaded. For more information, see [2.1.3 First Fault Register, FFR](#).

Packed access

A memory access that is performed as a result of a load or store instruction for which the vector element size and the memory element size are the same.

Predicate

A one-dimensional array of predicate elements of the same size.

Predicate element

Individual subdivisions of a predicate register that can be 1, 2, 4, or 8 bits in size. The predicate element size is specified independently by each instruction and is always one-eighth the size of the corresponding vector element. The lowest-numbered bit of each predicate element holds the Boolean value of that element, where 1 represents TRUE and 0 represents FALSE.

Predicate register

An *SVE* predicate register, P0-P15, having a length that is a multiple of 16 bits, in the range 16 to 256, inclusive.

Predicated instruction

When the instruction specifies a Governing predicate register.

Prefixed instruction

The instruction that immediately follows a MOVPRFX instruction in program order.

Scalar base register

A scalar base register refers to an AArch64 *general-purpose register*, X0-X31, or the current stack pointer, SP.

Scalar index register

A scalar index register refers to an AArch64 *general-purpose register*, X0-X31, or for certain instructions, XZR.

Scatter-store

Scatter-store is a mechanism that allows the elements of a vector to be written to non-contiguous memory locations using a vector of addresses, where the addresses are constructed according to the addressing mode. For more information, see [11.9 Load/store/prefetch instructions](#).

SIMD

Single Instruction, Multiple Data. A *SIMD* instruction performs the same operation on multiple vector elements or predicate elements in parallel.

Unpacked access

A memory access that is performed as a result of a load or store instruction for which the vector element size is larger than the memory element size.

Vector

A one-dimensional array of vector elements of the same size and data type.

Vector element

Individual subdivisions of a vector register that can be 8, 16, 32, 64 or 128 bits in size. The vector element size and data type is specified independently by each instruction.

Vector length

The accessible width of the *SVE* vector registers at the current *Exception level*, as constrained by the *ZCR_EL1*, *ZCR_EL2*, and *ZCR_EL3 System registers*. All vector registers at the same *Exception level* have the same vector length. The accessible width of the *SVE* predicate registers and *FFR* is one-eighth of the vector length.

Vector register

An *SVE* vector register, *Z0-Z31*, having a length that is a multiple of 128 bits, in the range 128 bits to 2048 bits, inclusive.

Zeroing predication

When a predicated instruction specifies zeroing predication, the Inactive elements of the destination register are set to zero.

Draft